

UNDERSTANDING TRAINABLE SPARSE CODING VIA MATRIX FACTORIZATION

Thomas Moreau

CMLA, ENS Cachan, CNRS,
 Université Paris-Saclay,
 94235 Cachan, France
 thomas.moreau@cmla.ens-cachan.fr

Joan Bruna

Courant Institute of Mathematical Sciences,*
 New York University,
 New York, NY 10012, USA
 joan.bruna@berkeley.edu

ABSTRACT

Sparse coding is a core building block in many data analysis and machine learning pipelines. Typically it is solved by relying on generic optimization techniques, such as the Iterative Soft Thresholding Algorithm and its accelerated version (ISTA, FISTA). These methods are optimal in the class of first-order methods for non-smooth, convex functions. However, they do not exploit the particular structure of the problem at hand nor the input data distribution. An acceleration using neural networks, coined LISTA, was proposed in [Gregor & Le Cun \(2010\)](#), which showed empirically that one could achieve high quality estimates with few iterations by modifying the parameters of the proximal splitting appropriately.

In this paper we study the reasons for such acceleration. Our mathematical analysis reveals that it is related to a specific matrix factorization of the Gram kernel of the dictionary, which attempts to nearly diagonalise the kernel with a basis that produces a small perturbation of the ℓ_1 ball. When this factorization succeeds, we prove that the resulting splitting algorithm enjoys an improved convergence bound with respect to the non-adaptive version. Moreover, our analysis also shows that conditions for acceleration occur mostly at the beginning of the iterative process, consistent with numerical experiments. We further validate our analysis by showing that on dictionaries where this factorization does not exist, adaptive acceleration fails.

1 INTRODUCTION

Feature selection is a crucial point in high dimensional data analysis. Different techniques have been developed to tackle this problem efficiently, and amongst them sparsity has emerged as a leading paradigm. In statistics, the LASSO estimator ([Tibshirani, 1996](#)) provides a reliable way to select features and has been extensively studied in the last two decades ([Hastie et al. \(2015\)](#) and references therein). In machine learning and signal processing, sparse coding has made its way into several modern architectures, including large scale computer vision ([Coates & Ng, 2011](#)) and biologically inspired models ([Cadiou & Olshausen, 2012](#)). Also, Dictionary learning is a generic unsupervised learning method to perform non-linear dimensionality reduction with efficient computational complexity ([Mairal et al., 2009](#)). All these techniques heavily rely on the resolution of ℓ_1 -regularized least squares.

The ℓ_1 -sparse coding problem is defined as solving, for a given input $x \in \mathbb{R}^n$ and dictionary $D \in \mathbb{R}^{n \times m}$, the following problem:

$$z^*(x) = \arg \min_z F_x(z) \triangleq \frac{1}{2} \|x - Dz\|^2 + \lambda \|z\|_1 . \quad (1)$$

This problem is convex and can therefore be solved using convex optimization machinery. Proximal splitting methods ([Beck & Teboulle, 2009](#)) alternate between the minimization of the smooth and differentiable part using the gradient information and the minimization of the non-differentiable part using a proximal operator ([Combettes & Bauschke, 2011](#)). These methods can also be accelerated by considering a momentum term, as it is done in FISTA

*Work done while appointed at UC Berkeley, Statistics Department (currently on leave)

(Beck & Teboulle, 2009; Nesterov, 2005). Coordinate descent (Friedman et al., 2007; Osher & Li, 2009) leverages the closed formula that can be derived for optimizing the problem (1) for one coordinate z_i given that all the other are fixed. At each step of the algorithm, one coordinate is updated to its optimal value, which yields an inexpensive scheme to perform each step. The choice of the coordinate to update at each step is critical for the performance of the optimization procedure. Least Angle Regression (LARS) (Hesterberg et al., 2008) is another method that computes the whole LASSO regularization path. These algorithms all provide an optimization procedure that leverages the local properties of the cost function iteratively. They can be shown to be optimal among the class of first-order methods for generic convex, non-smooth functions (Bubeck, 2014).

But all these results are given in the worst case and do not use the distribution of the considered problem. One can thus wonder whether a more efficient algorithm to solve (1) exists for a fixed dictionary D and generic input x drawn from a certain input data distribution. In Gregor & Le Cun (2010), the authors introduced LISTA, a trained version of ISTA that adapts the parameters of the proximal splitting algorithm to approximate the solution of the LASSO using a finite number of steps. This method exploits the common structure of the problem to learn a better transform than the generic ISTA step. As ISTA is composed of a succession of linear operations and piecewise non linearities, the authors use the neural network framework and the backpropagation to derive an efficient procedure solving the LASSO problem. In Sprechmann et al. (2012), the authors extended LISTA to more generic sparse coding scenarios and showed that adaptive acceleration is possible under general input distributions and sparsity conditions.

In this paper, we are interested in the following question: Given a finite computational budget, what is the optimum estimator of the sparse coding? This question belongs to the general topic of computational tradeoffs in statistical inference. Randomized sketches (Alaoui & Mahoney, 2015; Yang et al., 2015) reduce the size of convex problems by projecting expensive kernel operators into random subspaces, and reveal a tradeoff between computational efficiency and statistical accuracy. Agarwal (2012) provides several theoretical results on performing inference under various computational constraints, and Chandrasekaran & Jordan (2013) considers a hierarchy of convex relaxations that provide practical tradeoffs between accuracy and computational cost. More recently, Oymak et al. (2015) provides sharp time-data tradeoffs in the context of linear inverse problems, showing the existence of a phase transition between the number of measurements and the convergence rate of the resulting recovery optimization algorithm. Giryes et al. (2016) builds on this result to produce an analysis of LISTA that describes acceleration in conditions where the iterative procedure has linear convergence rate. Finally, Xin et al. (2016) also studies the capabilities of Deep Neural networks at approximating sparse inference. The authors show that unrolled iterations lead to better approximation if one allows the weights to vary at each layer, contrary to standard splitting algorithms. Whereas their focus is on relaxing the convergence hypothesis of iterative thresholding algorithms, we study a complementary question, namely when is speedup possible, without assuming strongly convex optimization. Their results are consistent with ours, since our analysis also shows that learning shared layer weights is less effective.

Inspired by the LISTA architecture, our mathematical analysis reveals that adaptive acceleration is related to a specific matrix factorization of the Gram matrix of the dictionary $B = D^T D$ as $B = A^T S A - R$, where A is unitary, S is diagonal and the residual is positive semidefinite: $R \succeq 0$. Our factorization balances between near diagonalization by asking that $\|R\|$ is small and small perturbation of the ℓ_1 norm, *i.e.* $\|Az\|_1 - \|z\|_1$ is small. When this factorization succeeds, we prove that the resulting splitting algorithm enjoys a convergence rate with improved constants with respect to the non-adaptive version. Moreover, our analysis also shows that acceleration is mostly possible at the beginning of the iterative process, when the current estimate is far from the optimal solution, which is consistent with numerical experiments. We also show that the existence of this factorization is not only sufficient for acceleration, but also necessary. This is shown by constructing dictionaries whose Gram matrix diagonalizes in a basis that is incoherent with the canonical basis, and verifying that LISTA fails in that case to accelerate with respect to ISTA.

In our numerical experiments, we design a specialized version of LISTA called FacNet, with more constrained parameters, which is then used as a tool to show that our theoretical analysis captures the acceleration mechanism of LISTA. Our theoretical results can be applied to FacNet and as LISTA is a generalization of this model, it always performs at least as well, showing that the existence of the factorization is a sufficient certificate for acceleration by

LISTA. Reciprocally, we show that for cases where no acceleration is possible with FacNet, the LISTA model also fail to provide acceleration, linking the two speedup mechanisms. This numerical evidence suggest that the existence of our proposed factorization is sufficient and somewhat necessary for LISTA to show good results.

The rest of the paper is structured as follows. [Section 2](#) presents our mathematical analysis and proves the convergence of the adaptive algorithm as a function of the quality of the matrix factorization. Finally, [Section 3](#) presents the generic architectures that will enable the usage of such schemes and the numerical experiments, which validate our analysis over a range of different scenarios.

2 ACCELERATING SPARSE CODING WITH SPARSE MATRIX FACTORIZATIONS

2.1 UNITARY PROXIMAL SPLITTING

In this section we describe our setup for accelerating sparse coding based on the Proximal Splitting method. Let $\Omega \subset \mathbb{R}^n$ be the set describing our input data, and $D \in \mathbb{R}^{n \times m}$ be a dictionary, with $m > n$. We wish to find fast and accurate approximations of the sparse coding $z^*(x)$ of any $x \in \Omega$, defined in [\(1\)](#) For simplicity, we denote $B = D^\top D$ and $y = D^\dagger x$ to rewrite [\(1\)](#) as

$$z^*(x) = \arg \min_z F_x(z) = \underbrace{\frac{1}{2}(y-z)^\top B(y-z)}_{E(z)} + \underbrace{\lambda \|z\|_1}_{G(z)}. \quad (2)$$

For clarity, we will refer to F_x as F and to $z^*(x)$ as z^* . The classic proximal splitting technique finds z^* as the limit of sequence $(z_k)_k$, obtained by successively constructing a surrogate loss $F_k(z)$ of the form

$$F_k(z) = E(z_k) + (z_k - y)^\top B(z - z_k) + L_k \|z - z_k\|_2^2 + \lambda \|z\|_1, \quad (3)$$

satisfying $F_k(z) \geq F(z)$ for all $z \in \mathbb{R}^m$. Since F_k is separable in each coordinate of z , $z_{k+1} = \arg \min_z F_k(z)$ can be computed efficiently. This scheme is based on a majoration of the quadratic form $(y-z)^\top B(y-z)$ with an isotropic quadratic form $L_k \|z_k - z\|_2^2$. The convergence rate of the splitting algorithm is optimized by choosing L_k as the smallest constant satisfying $F_k(z) \geq F(z)$, which corresponds to the largest singular value of B .

The computation of z_{k+1} remains separable by replacing the quadratic form $L_k \mathbf{I}$ by any diagonal form. However, the Gram matrix $B = D^\top D$ might be poorly approximated via diagonal forms for general dictionaries. Our objective is to accelerate the convergence of this algorithm by finding appropriate factorizations of the matrix B such that

$$B \approx A^\top S A, \quad \text{and} \quad \|Az\|_1 \approx \|z\|_1,$$

where A is unitary and S is diagonal positive definite. Given a point z_k at iteration k , we can rewrite $F(z)$ as

$$F(z) = E(z_k) + (z_k - y)^\top B(z - z_k) + Q_B(z, z_k), \quad (4)$$

with $Q_B(v, w) := \frac{1}{2}(v - w)^\top B(v - w) + \lambda \|v\|_1$. For any diagonal positive definite matrix S and unitary matrix A , the surrogate loss $\tilde{F}(z, z_k) := E(z_k) + (z_k - y)^\top B(z - z_k) + Q_S(Az, Az_k)$ can be explicitly minimized, since

$$\begin{aligned} \arg \min_z \tilde{F}(z, z_k) &= A^\top \arg \min_u \left((z_k - y)^\top B A^\top (u - Az_k) + Q_S(u, Az_k) \right) \\ &= A^\top \arg \min_u Q_S \left(u, Az_k - S^{-1} A B (z_k - y) \right) \end{aligned} \quad (5)$$

where we use the variable change $u = Az$. As S is diagonal positive definite, [\(5\)](#) is separable and can be computed easily, using a linear operation followed by a point-wise non linear soft-thresholding. Thus, any couple (A, S) ensures an computationally cheap scheme. The question is then how to factorize B using S and A in an optimal manner, that is, such that the resulting proximal splitting sequence converges as fast as possible to the sparse coding solution.

2.2 NON-ASYMPTOTIC ANALYSIS

We will now establish convergence results based on the previous factorization. These bounds will inform us on how to best choose the factors A_k and S_k in each iteration.

For that purpose, let us define

$$\delta_A(z) = \lambda \left(\|Az\|_1 - \|z\|_1 \right), \text{ and } R = A^\top SA - B. \quad (6)$$

The quantity $\delta_A(z)$ thus measures how invariant the ℓ_1 norm is to the unitary operator A , whereas R corresponds to the residual of approximating the original Gram matrix B by our factorization $A^\top SA$. Given a current estimate z_k , we can rewrite

$$\tilde{F}(z, z_k) = F(z) + \frac{1}{2}(z - z_k)^\top R(z - z_k) + \delta_A(z). \quad (7)$$

By imposing that R is a positive semidefinite residual one immediately obtains the following bound.

Proposition 2.1. *Suppose that $R = A^\top SA - B$ is positive definite, and define*

$$z_{k+1} = \arg \min_z \tilde{F}(z, z_k). \quad (8)$$

$$\text{Then } F(z_{k+1}) - F(z^*) \leq \frac{1}{2} \|R\| \|z_k - z^*\|_2^2 + \delta_A(z^*) - \delta_A(z_{k+1}). \quad (9)$$

Proof. By definition of z_{k+1} and using the fact that $R \succ 0$ we have

$$\begin{aligned} F(z_{k+1}) - F(z^*) &\leq F(z_{k+1}) - \tilde{F}(z_{k+1}, z_k) + \tilde{F}(z^*, z_k) - F(z^*) \\ &= -\frac{1}{2}(z_{k+1} - z_k)^\top R(z_{k+1} - z_k) - \delta_A(z_{k+1}) + \frac{1}{2}(z^* - z_k)^\top R(z^* - z_k) + \delta_A(z^*) \\ &\leq \frac{1}{2}(z^* - z_k)^\top R(z^* - z_k) + \left(\delta_A(z^*) - \delta_A(z_{k+1}) \right). \end{aligned}$$

where the first line results from the definition of z_{k+1} and the third line makes use of R positiveness. \square

This simple bound reveals that to obtain fast approximations to the sparse coding it is sufficient to find S and A such that $\|R\|$ is small and that the ℓ_1 commutation term δ_A is small. These two conditions will be often in tension: one can always obtain $R \equiv 0$ by using the Singular Value Decomposition of $B = A_0^\top S_0 A_0$ and setting $A = A_0$ and $S = S_0$. However, the resulting A_0 might introduce large commutation error δ_{A_0} . Similarly, as the absolute value is non-expansive, *i.e.* $|a| - |b| \leq |a - b|$, we have that

$$\begin{aligned} |\delta_A(z)| = \lambda \left| \|Az\|_1 - \|z\|_1 \right| &\leq \lambda \|(A - \mathbf{I})z\|_1 \\ &\leq \lambda \sqrt{2 \max(\|Az\|_0, \|z\|_0)} \cdot \|A - \mathbf{I}\| \cdot \|z\|_2, \end{aligned} \quad (10)$$

where we have used the Cauchy-Schwartz inequality $\|x\|_1 \leq \sqrt{\|x\|_0} \|x\|_2$ in the last equation. In particular, (10) shows that unitary matrices in the neighborhood of \mathbf{I} with $\|A - \mathbf{I}\|$ small have small ℓ_1 commutation error δ_A but can be inappropriate to approximate general B matrix.

The commutation error also depends upon the sparsity of z and Az . If both z and Az are sparse then the commutation error is reduced, which can be achieved if A is itself a sparse unitary matrix. Moreover, since

$$|\delta_A(z) - \delta_A(z')| \leq \lambda \|z\|_1 - \|z'\|_1 + \lambda \|Az\|_1 - \|Az'\|_1$$

$$\text{and } \|z\|_1 - \|z'\|_1 \leq \|z - z'\|_1 \leq \sqrt{\|z - z'\|_0} \|z - z'\|_2$$

it results that δ_A is Lipschitz with respect to the Euclidean norm; let us denote by $L_A(z)$ its local Lipschitz constant in z , which can be computed using the norm of the subgradient

in z^1 . An uniform upper bound for this constant is $(1 + \|A\|_1)\lambda\sqrt{m}$, but it is typically much smaller when z and Az are both sparse.

Equation (8) defines an iterative procedure determined by the pairs $\{(A_k, S_k)\}_k$. The following theorem uses the previous results to compute an upper bound of the resulting sparse coding estimator.

Theorem 2.2. *Let A_k, S_k be the pair of unitary and diagonal matrices corresponding to iteration k , chosen such that $R_k = A_k^\top S_k A_k - B \succ 0$. It results that*

$$F(z_k) - F(z^*) \leq \frac{(z^* - z_0)^\top R_0 (z^* - z_0) + 2L_{A_0}(z_1)\|z^* - z_1\|_2}{2k} + \frac{\alpha - \beta}{2k}, \quad (11)$$

$$\begin{aligned} \text{with } \alpha &= \sum_{i=1}^{k-1} \left(2L_{A_i}(z_{i+1})\|z^* - z_{i+1}\|_2 + (z^* - z_i)^\top (R_{i-1} - R_i)(z^* - z_i) \right), \\ \beta &= \sum_{i=0}^{k-1} (i+1) \left((z_{i+1} - z_i)^\top R_i (z_{i+1} - z_i) + 2\delta_{A_i}(z_{i+1}) - 2\delta_{A_i}(z_i) \right), \end{aligned}$$

where $L_A(z)$ denote the local lipschitz constant of δ_A at z .

Remarks: If one sets $A_k = \mathbf{I}$ and $S_k = \|B\|\mathbf{I}$ for all $k \geq 0$, (11) corresponds to the bound of the ISTA algorithm (Beck & Teboulle, 2009).

We can specialize the theorem in the case when A_0, S_0 are chosen to minimize the bound (9) and $A_k = \mathbf{I}$, $S_k = \|B\|\mathbf{I}$ for $k \geq 1$.

Corollary 2.3. *If $A_k = \mathbf{I}$, $S_k = \|B\|\mathbf{I}$ for $k \geq 1$ then*

$$F(z_k) - F(z^*) \leq \frac{(z^* - z_0)^\top R_0 (z^* - z_0) + 2L_{A_0}(z_1)(\|z^* - z_1\| + \|z_1 - z_0\|) + (z^* - z_1)^\top R_0 (z^* - z_1)^\top}{2k}. \quad (12)$$

This corollary shows that by simply replacing the first step of ISTA by the modified proximal step detailed in (5), one can obtain an improved bound at fixed k as soon as

$$2\|R_0\| \max(\|z^* - z_0\|_2^2, \|z^* - z_1\|_2^2) + 4L_{A_0}(z_1) \max(\|z^* - z_0\|_2, \|z^* - z_1\|_2) \leq \|B\| \|z^* - z_0\|_2^2,$$

which, assuming $\|z^* - z_0\|_2 \geq \|z^* - z_1\|_2$, translates into

$$\|R_0\| + 2 \frac{L_{A_0}(z_1)}{\|z^* - z_0\|_2} \leq \frac{\|B\|}{2}. \quad (13)$$

More generally, given a current estimate z_k , searching for a factorization (A_k, S_k) will improve the upper bound when

$$\|R_k\| + 2 \frac{L_{A_k}(z_{k+1})}{\|z^* - z_k\|_2} \leq \frac{\|B\|}{2}. \quad (14)$$

We emphasize that this is not a guarantee of acceleration, since it is based on improving an upper bound. However, it provides a simple picture on the mechanism that makes non-asymptotic acceleration possible.

2.3 INTERPRETATION

In this section we analyze the consequences of Theorem 2.2 in the design of fast sparse coding approximations, and provide a possible explanation for the behavior observed numerically.

2.3.1 ‘PHASE TRANSITION’ AND LAW OF DIMINISHING RETURNS

(14) reveals that the optimum matrix factorization in terms of minimizing the upper bound depends upon the current scale of the problem, that is, of the distance $\|z^* - z_k\|$. At the beginning of the optimization, when $\|z^* - z_k\|$ is large, the bound (14) makes it easier to explore the space of factorizations (A, S) with A further away from the identity. Indeed, the bound tolerates larger increases in $L_A(z_{k+1})$, which is dominated by

$$L_A(z_{k+1}) \leq \lambda(\sqrt{\|z_{k+1}\|_0} + \sqrt{\|Az_{k+1}\|_0}),$$

¹ This quantity exists as δ_A is a difference of convex. See proof of ?? in appendices for precisions.




Figure 1: Network architecture for ISTA/LISTA. The unfolded version (b) is trainable through backpropagation and permits to approximate the sparse coding solution efficiently.

i.e. the sparsity of both z_1 and $A_0(z_1)$. On the other hand, when we reach intermediate solutions z_k such that $\|z^* - z_k\|$ is small with respect to $L_A(z_{k+1})$, the upper bound is minimized by choosing factorizations where A is closer and closer to the identity, leading to the non-adaptive regime of standard ISTA ($A = Id$).

This is consistent with the numerical experiments, which show that the gains provided by learned sparse coding methods are mostly concentrated in the first iterations. Once the estimates reach a certain energy level, section 3 shows that LISTA enters a steady state in which the convergence rate matches that of standard ISTA.

The natural follow-up question is to determine how many layers of adaptive splitting are sufficient before entering the steady regime of convergence. A conservative estimate of this quantity would require an upper bound of $\|z^* - z_k\|$ from the energy bound $F(z_k) - F(z^*)$. Since in general F is convex but not strongly convex, such bound does not exist unless one can assume that F is locally strongly convex (for instance for sufficiently small values of F).

2.3.2 IMPROVING THE FACTORIZATION TO PARTICULAR INPUT DISTRIBUTIONS

Given an input dataset $\mathcal{D} = (x_i, z_i^{(0)}, z_i^*)_{i \leq N}$, containing examples $x_i \in \mathbb{R}^n$, initial estimates $z_i^{(0)}$ and sparse coding solutions z_i^* , the factorization adapted to \mathcal{D} is defined as

$$\min_{A, S; A^T A = \mathbf{I}, A^T S A - B > 0} \frac{1}{N} \sum_{i \leq N} \frac{1}{2} (z_i^{(0)} - z_i^*)^T (A^T S A - B) (z_i^{(0)} - z_i^*) + \delta_A(z_i^*) - \delta_A(z_{1,i}) . \quad (15)$$

Therefore, adapting the factorization to a particular dataset, as opposed to enforcing it uniformly over a given ball $B(z^*; R)$ (where the radius R ensures that the initial value $z_0 \in B(z^*; R)$), will always improve the upper bound (9). Studying the gains resulting from the adaptation to the input distribution will be let for future work.

3 NUMERICAL EXPERIMENTS

This section provides numerical arguments to analyse adaptive optimization algorithms and their performances, and relates them to the theoretical properties developed in the previous section. All the experiments were run using Python and Tensorflow. For all the experiments, the training is performed using Adagrad (Duchi et al., 2011). The code to reproduce the figures is available online².

3.1 ADAPTIVE OPTIMIZATION NETWORKS ARCHITECTURES

LISTA/LFISTA In Gregor & Le Cun (2010), the authors introduced LISTA, a neural network constructed by considering ISTA as a recurrent neural net. At each step, ISTA performs the following 2-step procedure :

$$\left. \begin{aligned} 1. \quad u_{k+1} &= z_k - \frac{1}{L} D^T (D z_k - x) = \underbrace{(\mathbf{I} - \frac{1}{L} D^T D)}_{W_g} z_k + \underbrace{\frac{1}{L} D^T x}_{W_e} , \\ 2. \quad z_{k+1} &= h_{\frac{\lambda}{L}}(u_{k+1}) \text{ where } h_{\theta}(u) = \text{sign}(u)(|u| - \theta)_+ , \end{aligned} \right\} \text{step } k \text{ of ISTA} \quad (16)$$

²The code can be found at <https://github.com/tomMoral/AdaptiveOptim>

This procedure combines a linear operation to compute u_{k+1} with an element-wise non linearity. It can be summarized as a recurrent neural network, presented in [Figure 1a.](#), with tied weights. The authors in [Gregor & Le Cun \(2010\)](#) considered the architecture Φ_{Θ}^K with parameters $\Theta = (W_g^{(k)}, W_e^{(k)}, \theta^{(k)})_{k=1, \dots, K}$ obtained by unfolding K times the recurrent network, as presented in [Figure 1b.](#) The layers ϕ_{Θ}^k are defined as

$$z_{k+1} = \phi_{\Theta}^k(z_k) := h_{\theta}(W_g z_k + W_e x) . \quad (17)$$

If $W_g^{(k)} = \mathbf{I} - \frac{D^T D}{L}$, $W_e^{(k)} = \frac{D^T}{L}$ and $\theta^{(k)} = \frac{\lambda}{L}$ are fixed for all the K layers, the output of this neural net is exactly the vector z_K resulting from K steps of ISTA. With LISTA, the parameters Θ are learned using back propagation to minimize the cost function: $f(\Theta) = \mathbb{E}_x \left[F_x(\Phi_{\Theta}^K(x)) \right]$.

A similar algorithm can be derived from FISTA, the accelerated version of ISTA to obtain LFISTA (see [Figure 5](#) in [Appendix A](#)). The architecture is very similar to LISTA, now with two memory tapes:

$$z_{k+1} = h_{\theta}(W_g z_k + W_m z_{k-1} + W_e x) .$$

Factorization network Our analysis in [Section 2](#) suggests a refactorization of LISTA in more a structured class of parameters. Following the same basic architecture, and using [\(5\)](#), the network FacNet, Ψ_{Θ}^K is formed using layers such that:

$$z_{k+1} = \psi_{\Theta}^k(z_k) := A^T h_{\lambda S^{-1}}(A z_k - S^{-1} A (D^T D z_k - D^T x)) , \quad (18)$$

with S diagonal and A unitary, the parameters of the k -th layer. The parameters obtained after training such a network with back-propagation can be used with the theory developed in [Section 2](#). Up to the last linear operation A^T of the network, this network is a re-parametrization of LISTA in a more constrained parameter space. Thus, LISTA is a generalization of this proposed network and should have performances at least as good as FacNet, for a fixed number of layers.

The optimization can also be performed using backpropagation. To enforce the unitary constraints on $A^{(k)}$, the cost function is modified with a penalty:

$$f(\Theta) = \mathbb{E}_x \left[F_x(\Psi_{\Theta}^K(x)) \right] + \frac{\mu}{K} \sum_{k=1}^K \left\| \mathbf{I} - \left(A^{(k)} \right)^T A^{(k)} \right\|_2^2 , \quad (19)$$

with $\Theta = (A^{(k)}, S^{(k)})_{k=1 \dots K}$ the parameters of the K layers and μ a scaling factor for the regularization. The resulting matrix $A^{(k)}$ is then projected on the Stiefel Manifold using a SVD to obtain final parameters, coherent with the network structure.

Linear model Finally, it is important to distinguish the performance gain resulting from choosing a suitable starting point and the acceleration from our model. To highlights the gain obtain by changing the starting point, we considered a linear model with one layer such that $z_{out} = A^{(0)} x$. This model is learned using SGD with the convex cost function $f(A^{(0)}) = \|\mathbf{I} - D A^{(0)}\|_2^2 + \lambda \|A^{(0)} x\|_1$. It computes a tradeoff between starting from the sparsest point $\mathbf{0}$ and a point with minimal reconstruction error y . Then, we observe the performance of the classical iteration of ISTA using z_{out} as a starting point instead of $\mathbf{0}$.

3.2 SYNTHETIC PROBLEMS WITH KNOWN DISTRIBUTIONS

Gaussian dictionary In order to disentangle the role of dictionary structure from the role of data distribution structure, the minimization problem is tested using a synthetic generative model with no structure in the weights distribution. First, m atoms $d_i \in \mathbb{R}^n$ are drawn *iid* from a multivariate Gaussian with mean $\mathbf{0}$ and covariance \mathbf{I}_n and the dictionary D is defined as $\left(d_i / \|d_i\|_2 \right)_{i=1 \dots m}$. The data points are generated from its sparse codes following a Bernoulli-Gaussian model. The coefficients $z = (z_1, \dots, z_m)$ are constructed with $z_i = b_i a_i$, where $b_i \sim \mathcal{B}(\rho)$ and $a_i \sim \mathcal{N}(0, \sigma \mathbf{I}_m)$, where ρ controls the sparsity of the data. The values are set to $m=100$, $n=64$ for the dictionary dimension, $\rho = 5/m$ for the sparsity level and $\sigma=10$ for the activation coefficient generation parameters. The sparsity




Figure 2: Evolution of the cost function $F(z_k) - F(z^*)$ with the number of layers or the number of iteration k for different sparsity level. (left) $\rho = 1/20$ and (right) $\rho = 1/4$.

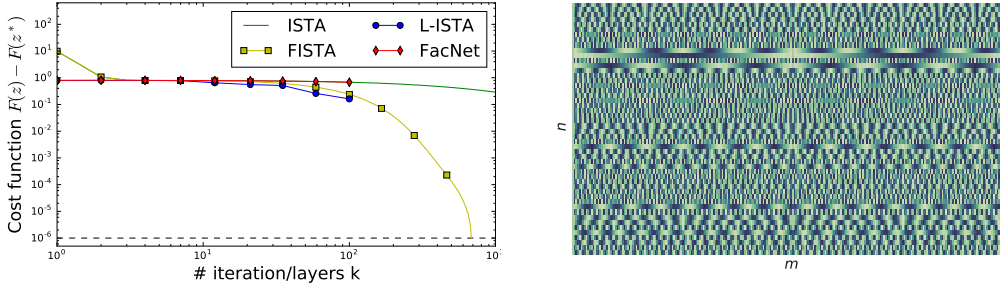


Figure 3: Evolution of the cost function $F(z_k) - F(z^*)$ with the number of layers or the number of iteration k for a problem generated with an adversarial dictionary.

regularization is set to $\lambda=0.01$. The batches used for the training are generated with the model at each step and the cost function is evaluated over a fixed test set, not used in the training.

Figure 2 displays the cost performance for methods ISTA/FISTA/Linear relatively to their iterations and for methods LISTA/LFISTA/FacNet relatively to the number of layers used to solve our generated problem. Linear has performances comparable to learned methods with the first iteration but a gap appears as the number of layers increases, until a point where it achieves the same performances as non adaptive methods. This highlights that the adaptation is possible in the subsequent layers of the networks, going farther than choosing a suitable starting point for iterative methods. The first layers permit to achieve a large gain over the classical optimization strategy, by leveraging the structure of the problem. This appears even with no structure in the sparsity patterns of input data, in accordance with the results in the previous section. We also observe diminishing returns as the number of layers increases. This results from the phase transition described in Subsubsection 2.3.1, as the last layers behave as ISTA steps and do not speed up the convergence. The 3 learned algorithms are always performing at least as well as their classical counterpart, as it was stated in Theorem 2.2. We also explored the effect of the sparsity level in the training and learning of adaptive networks. In the denser setting, the arbitrage between the ℓ_1 -norm and the squared error is easier as the solution has a lot of non zero coefficients. Thus in this setting, the approximate method is more precise than in the very sparse setting where the approximation must perform a fine selection of the coefficients. But it also yield lower gain at the beginning as the sparser solution can move faster.

There is a small gap between LISTA and FacNet in this setup. This can be explained from the extra constraints on the weights that we impose in the FacNet, which effectively reduce the parameter space by half. Also, we implement the unitary constraints on the matrix A by a soft regularization (see (19)), involving an extra hyper-parameter μ that also contributes to the small performance gap. In any case, these experiments show that our analysis accounts for most of the acceleration provided by LISTA, as the performance of both methods are similar, up to optimization errors.

Adversarial dictionary The results from Section 2 show that problems with a gram matrix composed of large eigenvalues associated to non sparse eigenvectors are harder to accelerate. Indeed, it is not possible in this case to find a quasi diagonalization of the matrix B that

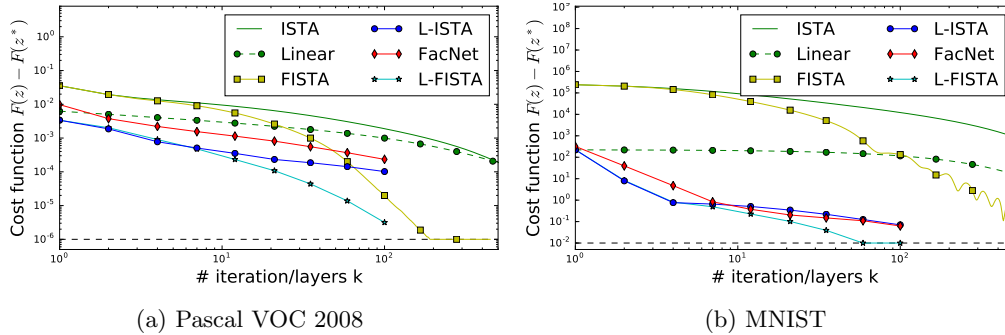


Figure 4: Evolution of the cost function $F(z_k) - F(z^*)$ with the number of layers or the number of iteration k for two image datasets.

does not distort the ℓ_1 norm. It is possible to generate such a dictionary using Harmonic Analysis. The Discrete Fourier Transform (DFT) distorts a lot the ℓ_1 ball, since a very sparse vector in the temporal space is transformed in widely spread spectrum in the Fourier domain. We can thus design a dictionary for which LISTA and FacNet performances should be degraded. $D = \left(d_i / \|d_i\|_2 \right)_{i=1 \dots m}$ is constructed such that $d_{j,k} = e^{-2\pi i j \zeta_k}$, with $(\zeta_k)_{k \leq n}$ randomly selected from $\left\{ 1/m, \dots, m/2/m \right\}$ without replacement.

The resulting performances are reported in [Figure 3](#). The first layer provides a big gain by changing the starting point of the iterative methods. It realizes an arbitrage of the tradeoff between starting from $\mathbf{0}$ and starting from y . But the next layers do not yield any extra gain compared to the original ISTA algorithm. After 4 layers, the cost performance of both adaptive methods and ISTA are equivalent. It is clear that in this case, FacNet does not accelerate efficiently the sparse coding, in accordance with our result from [Section 2](#). LISTA also displays poor performances in this setting. This provides further evidence that FacNet and LISTA share the same acceleration mechanism as adversarial dictionaries for FacNet are also adversarial for LISTA.

3.3 SPARSE CODING WITH OVER COMPLETE DICTIONARY ON IMAGES

Wavelet encoding for natural images A highly structured dictionary composed of translation invariant Haar wavelets is used to encode 8×8 patches of images from the PASCAL VOC 2008 dataset. The network is used to learn an efficient sparse coder for natural images over this family. 500 images are sampled from dataset to train the encoder. Training batches are obtained by uniformly sampling patches from the training image set to feed the stochastic optimization of the network. The encoder is then tested with 10000 patches sampled from 100 new images from the same dataset.

Learned dictionary for MNIST To evaluate the performance of LISTA for dictionary learning, LISTA was used to encode MNIST images over an unconstrained dictionary, learned *a priori* using classical dictionary learning techniques. The dictionary of 100 atoms was learned from 10000 MNIST images in grayscale rescaled to 17×17 using the implementation of [Mairal et al. \(2009\)](#) proposed in scikit-learn, with $\lambda = 0.05$. Then, the networks were trained through backpropagation using all the 60000 images from the training set of MNIST. Finally, the performance of these encoders were evaluated with the 10000 images of the training set of MNIST.

The [Figure 4](#) displays the cost performance of the adaptive procedures compared to non-adaptive algorithms. In both scenario, FacNet has performances comparable to the one of LISTA and their behavior are in accordance with the theory developed in [Section 2](#). The gains become smaller for each added layer and the initial gain is achieved for dictionary either structured or unstructured. The MNIST case presents a much larger gain compare to the experiment with natural images. This results from the difference of structure of the input distribution, as the MNIST digits are much more constrained than patches from natural images and the network is able to leverage it to find a better encoder. In the MNIST case, a network composed of 12 layers is sufficient to achieve performance comparable to ISTA with more than 1000 iterations.

4 CONCLUSIONS

In this paper we studied the problem of finite computational budget approximation of sparse coding. Inspired by the ability of neural networks to accelerate over splitting methods on the first few iterations, we have studied which properties of the dictionary matrix and the data distribution lead to such acceleration. Our analysis reveals that one can obtain acceleration by finding approximate matrix factorizations of the dictionary which nearly diagonalize its Gram matrix, but whose orthogonal transformations leave approximately invariant the ℓ_1 ball. By appropriately balancing these two conditions, we show that the resulting rotated proximal splitting scheme has an upper bound which improves over the ISTA upper bound under appropriate sparsity.

In order to relate this specific factorization property to the actual LISTA algorithm, we have introduced a reparametrization of the neural network that specifically computes the factorization, and incidentally provides reduced learning complexity (less parameters) from the original LISTA. Numerical experiments of [Section 3](#) show that such reparametrization recovers the same gains as the original neural network, providing evidence that our theoretical analysis is partially explaining the behavior of the LISTA neural network. Our acceleration scheme is inherently transient, in the sense that once the iterates are sufficiently close to the optimum, the factorization is not effective anymore. This transient effect is also consistent with the performance observed numerically, although the possibility remains open to find alternative models that further exploit the particular structure of the sparse coding. Finally, we provide evidence that successful matrix factorization is not only sufficient but also necessary for acceleration, by showing that Fourier dictionaries are not accelerated.

Despite these initial results, a lot remains to be understood on the general question of optimal tradeoffs between computational budget and statistical accuracy. Our analysis so far did not take into account any probabilistic consideration (e.g. obtain approximations that hold with high probability or in expectation). Another area of further study is the extension of our analysis to the FISTA case, and more generally to other inference tasks that are currently solved via iterative procedures compatible with neural network parametrizations, such as inference in Graphical Models using Belief Propagation or other ill-posed inverse problems.

REFERENCES

- Alekh Agarwal. *Computational Trade-offs in Statistical Learning*. PhD thesis, University of California, Berkeley, 2012.
- Ahmed Alaoui and Michael W Mahoney. Fast randomized kernel ridge regression with statistical guarantees. In *Advances in Neural Information Processing Systems (NIPS)*, pp. 775–783, 2015.
- Amir Beck and Marc Teboulle. A Fast Iterative Shrinkage-Thresholding Algorithm for Linear Inverse Problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.
- Sébastien Bubeck. Theory of convex optimization for machine learning. *preprint*, arXiv:1405(4980), 2014.
- Charles F Cadieu and Bruno A Olshausen. Learning intermediate-level representations of form and motion from natural movies. *Neural computation*, 24(4):827–866, 2012.
- Venkat Chandrasekaran and Michael I Jordan. Computational and statistical tradeoffs via convex relaxation. *Proceedings of the National Academy of Sciences*, 110(13):E1181–E1190, 2013.
- Adam Coates and Andrew Y Ng. The importance of encoding versus training with sparse coding and vector quantization. In *Proceedings of the 28th International Conference on Machine Learning (ICML-11)*, pp. 921–928, 2011.
- Patrick L Combettes and Heinz H. Bauschke. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*, volume 1. 2011.
- John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *The Journal of Machine Learning Research*, 12: 2121–2159, 2011.

- Jerome Friedman, Trevor Hastie, Holger Höfling, and Robert Tibshirani. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007.
- Raja Giryes, Yonina C Eldar, Alex M Bronstein, and Guillermo Sapiro. Tradeoffs between convergence speed and reconstruction accuracy in inverse problems. *preprint*, arXiv:1605(09232), 2016.
- Karol Gregor and Yann Le Cun. Learning Fast Approximations of Sparse Coding. In *International Conference on Machine Learning (ICML)*, pp. 399–406, 2010.
- Trevor Hastie, Robert Tibshirani, and Martin J. Wainwright. *Statistical Learning with Sparsity*. CRC Press, 2015.
- Tim Hesterberg, Nam Hee Choi, Lukas Meier, and Chris Fraley. Least angle and 1 penalized regression: A review. *Statistics Surveys*, 2:61–93, 2008.
- J. B. Hiriart-Urruty. How to regularize a difference of convex functions. *Journal of Mathematical Analysis and Applications*, 162(1):196–209, 1991.
- Julien Mairal, Francis Bach, Jean Ponce, and Guillermo Sapiro. Online Learning for Matrix Factorization and Sparse Coding. *Journal of Machine Learning Research*, 11(1):19–60, 2009.
- Yu Nesterov. Smooth minimization of non-smooth functions. *Mathematical Programming*, 103(1):127–152, 2005.
- Stanley Osher and Yingying Li. Coordinate descent optimization for l1 minimization with application to compressed sensing; a greedy algorithm. *Inverse Problems and Imaging*, 3(3):487–503, 2009.
- Samet Oymak, Benjamin Recht, and Mahdi Soltanolkotabi. Sharp time–data tradeoffs for linear inverse problems. *preprint*, arXiv:1507(04793), 2015.
- Pablo Sprechmann, Alex Bronstein, and Guillermo Sapiro. Learning Efficient Structured Sparse Models. In *International Conference on Machine Learning (ICML)*, pp. 615–622, 2012.
- Robert Tibshirani. Regression Shrinkage and Selection via the Lasso. *Journal of the royal statistical society. Series B (methodological)*, 58(1):267–288, 1996.
- Bo Xin, Yizhou Wang, Wen Gao, and David Wipf. Maximal sparsity with deep networks? *preprint*, arXiv:1605(01636), 2016.
- Yun Yang, Mert Pilanci, and Martin J Wainwright. Randomized sketches for kernels: Fast and optimal non-parametric regression. *preprint*, arXiv:1501(06195), 2015.

A LEARNED FISTA

A similar algorithm can be derived from FISTA, the accelerated version of ISTA to obtain LFISTA (see [Figure 5](#)). The architecture is very similar to LISTA, now with two memory taps: It introduces a momentum term to improve the convergence rate of ISTA as follows:

1. $y_k = z_k + \frac{t_{k-1} - 1}{t_k} (z_k - z_{k-1})$,
2. $z_{k+1} = h_{\frac{\lambda}{L}} \left(y_k - \frac{1}{L} \nabla E(y_k) \right) = h_{\frac{\lambda}{L}} \left(\left(\mathbf{I} - \frac{1}{L} B \right) y_k + \frac{1}{L} D^\top x \right)$,
3. $t_{k+1} = \frac{1 + \sqrt{1 + 4t_k^2}}{2}$.

By substituting the expression for y_k into the first equation, we obtain a generic recurrent architecture very similar to LISTA, now with two memory taps, that we denote by LFISTA:

$$z_{k+1} = h_\theta (W_g^{(k)} z_k + W_m^{(k)} z_{k-1} + W_e^{(k)} x) .$$




Figure 5: Network architecture for LFISTA. This network is trainable through backpropagation and permits to approximate the sparse coding solution efficiently.

This model is equivalent to running K -steps of FISTA when its parameters are initialized with

$$\begin{aligned} W_g^{(k)} &= \left(1 + \frac{t_{k-1} - 1}{t_k}\right) \left(\mathbf{I} - \frac{1}{L}B\right), \\ W_m^{(k)} &= \left(\frac{1 - t_{k-1}}{t_k}\right) \left(\mathbf{I} - \frac{1}{L}B\right), \\ W_e^{(k)} &= \frac{1}{L}D^\top. \end{aligned}$$

The parameters of this new architecture, presented in [Figure 5](#), are trained analogously as in the LISTA case.

B PROOFS

Lemma B.1. *Suppose that $R = A^\top SA - B$ is positive definite, and define*

$$z_{k+1} = \arg \min_z \tilde{F}(z, z_k), \text{ and} \quad (20)$$

$\delta_A(z) = \|Az\|_1 - \|z\|_1$. Then we have

$$F(z_{k+1}) - F(z^*) \leq \frac{1}{2} \left((z^* - z_k)^\top R(z^* - z_k) - (z^* - z_{k+1})^\top R(z^* - z_{k+1}) \right) + \langle \partial\delta_A(z_{k+1}), z_{k+1} - z^* \rangle. \quad (21)$$

Proof. We define

$$f(t) = F\left(tz_{k+1} + (1-t)z^*\right), \quad t \in [0, 1].$$

Since F is convex, f is also convex in $[0, 1]$. Since $f(0) = F(z^*)$ is the global minimum, it results that $f'(t)$ is increasing in $(0, 1]$, and hence

$$F(z_{k+1}) - F(z^*) = f(1) - f(0) = \int_0^1 f'(t) dt \leq f'(1),$$

where $f'(1)$ is any element of $\partial f(1)$. Since $\delta_A(z)$ is a difference of convex functions, its subgradient can be defined as a limit of infimal convolutions [Hiriart-Urruty \(1991\)](#). We have

$$\partial f(1) = \langle \partial F(z_{k+1}), z_{k+1} - z^* \rangle,$$

and since

$$\partial F(z) = \partial \tilde{F}(z, z_k) - R(z - z_k) - \partial\delta_A(z) \quad \text{and} \quad 0 \in \partial \tilde{F}(z_{k+1}, z_k)$$

it results that

$$\partial F(z_{k+1}) = -R(z_{k+1} - z_k) - \partial\delta_A(z_{k+1}),$$

and thus

$$F(z_{k+1}) - F(z^*) \leq (z^* - z_{k+1})^\top R(z_{k+1} - z_k) + \langle \partial\delta_A(z_{k+1}), (z^* - z_{k+1}) \rangle. \quad (22)$$

(21) is obtained by observing that

$$(z^* - z_{k+1})^\top R(z_{k+1} - z_k) \leq \frac{1}{2} \left((z^* - z_k)^\top R(z^* - z_k) - (z^* - z_{k+1})^\top R(z^* - z_{k+1}) \right), \quad (23)$$

thanks to the fact that $R \succ 0$. \square

Theorem B.2. Let A_k, S_k be the pair of unitary and diagonal matrices corresponding to iteration k , chosen such that $R_k = A_k^\top S_k A_k - B \succ 0$. It results that

$$F(z_k) - F(z^*) \leq \frac{(z^* - z_0)^\top R_0 (z^* - z_0) + 2\langle \nabla \delta_{A_0}(z_1), (z^* - z_1) \rangle}{2k} + \frac{\alpha - \beta}{2k}, \quad \text{with} \quad (24)$$

$$\alpha = \sum_{n=1}^{k-1} \left(2\langle \nabla \delta_{A_n}(z_{n+1}), (z^* - z_{n+1}) \rangle + (z^* - z_n)^\top (R_{n-1} - R_n)(z^* - z_n) \right),$$

$$\beta = \sum_{n=0}^{k-1} (n+1) \left((z_{n+1} - z_n)^\top R_n (z_{n+1} - z_n) + 2\delta_{A_n}(z_{n+1}) - 2\delta_{A_n}(z_n) \right).$$

Proof: The proof is adapted from (Beck & Teboulle, 2009), Theorem 3.1. From Lemma B.1, we start by using (21) to bound terms of the form $F(z_n) - F(z^*)$:

$$F(z_n) - F(z^*) \leq \langle \nabla \delta_{A_n}(z_{n+1}), (z^* - z_{n+1}) \rangle + \frac{1}{2} \left((z^* - z_n)^\top R_n (z^* - z_n) - (z^* - z_{n+1})^\top R_n (z^* - z_{n+1}) \right).$$

Adding these inequalities for $n = 0 \dots k-1$ we obtain

$$\left(\sum_{n=0}^{k-1} F(z_n) \right) - kF(z^*) \leq \sum_{n=0}^{k-1} \langle \nabla \delta_{A_n}(z_{n+1}), (z^* - z_{n+1}) \rangle + \quad (25)$$

$$+ \frac{1}{2} \left((z^* - z_0)^\top R_0 (z^* - z_0) - (z^* - z_k)^\top R_{k-1} (z^* - z_k) \right) +$$

$$+ \frac{1}{2} \sum_{n=1}^{k-1} (z^* - z_n)^\top (R_{n-1} - R_n)(z^* - z_n).$$

On the other hand, we also have

$$F(z_n) - F(z_{n+1}) \geq F(z_n) - \tilde{F}(z_n, z_n) + \tilde{F}(z_{n+1}, z_n) - F(z_{n+1})$$

$$= -\delta_{A_n}(z_n) + \delta_{A_n}(z_{n+1}) + \frac{1}{2} (z_{n+1} - z_n)^\top R_n (z_{n+1} - z_n),$$

which results in

$$\sum_{n=0}^{k-1} (n+1) (F(z_n) - F(z_{n+1})) \geq \frac{1}{2} \sum_{n=0}^{k-1} (n+1) (z_{n+1} - z_n)^\top R_n (z_{n+1} - z_n) + \quad (26)$$

$$+ \sum_{n=0}^{k-1} (n+1) (\delta_{A_n}(z_{n+1}) - \delta_{A_n}(z_n))$$

$$\left(\sum_{n=0}^{k-1} F(z_n) \right) - kF(z_k) \geq \sum_{n=0}^{k-1} (n+1) \left(\frac{1}{2} (z_{n+1} - z_n)^\top R_n (z_{n+1} - z_n) + \delta_{A_n}(z_{n+1}) - \delta_{A_n}(z_n) \right).$$

Combining (25) and (26) we obtain

$$F(z_k) - F(z^*) \leq \frac{(z^* - z_0)^\top R_0 (z^* - z_0) + 2\langle \nabla \delta_{A_0}(z_1), (z^* - z_1) \rangle}{2k} + \frac{\alpha - \beta}{2k} \quad (27)$$

with

$$\alpha = \sum_{n=1}^{k-1} \left(2\langle \nabla \delta_{A_n}(z_{n+1}), (z^* - z_{n+1}) \rangle + (z^* - z_n)^\top (R_{n-1} - R_n)(z^* - z_n) \right),$$

$$\beta = \sum_{n=0}^{k-1} (n+1) \left((z_{n+1} - z_n)^\top R_n (z_{n+1} - z_n) + 2\delta_{A_n}(z_{n+1}) - 2\delta_{A_n}(z_n) \right).$$

□

Corollary B.3. If $A_k = \mathbf{I}$, $S_k = \|B\|\mathbf{I}$ for $k > 0$ then

$$F(z_k) - F(z^*) \leq \frac{(z^* - z_0)^\top R_0 (z^* - z_0) + 2L_{A_0}(z_1) (\|z^* - z_1\| + \|z_1 - z_0\|) + (z^* - z_1)^\top R_0 (z^* - z_1)^\top}{2k}. \quad (28)$$

Proof: We verify that in that case, $R_{n-1} - R_n \equiv 0$ and for $n > 1$ and $\delta_{A_n} \equiv 0$ for $n > 0$. □