



LSALSA: accelerated source separation via learned sparse coding

Benjamin Cowen¹ · Apoorva Nandini Saridena¹ · Anna Choromanska¹

Received: 22 October 2018 / Revised: 10 April 2019 / Accepted: 15 May 2019 / Published online: 21 May 2019
© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2019

Abstract

We propose an efficient algorithm for the generalized sparse coding (SC) inference problem. The proposed framework applies to both the single dictionary setting, where each data point is represented as a sparse combination of the columns of one dictionary matrix, as well as the multiple dictionary setting as given in morphological component analysis (MCA), where the goal is to separate a signal into additive parts such that each part has distinct sparse representation within an appropriately chosen corresponding dictionary. Both the SC task and its generalization via MCA have been cast as ℓ_1 -regularized optimization problems of minimizing quadratic reconstruction error. In an effort to accelerate traditional acquisition of sparse codes, we propose a deep learning architecture that constitutes a trainable time-unfolded version of the split augmented lagrangian shrinkage algorithm (SALSA), a special case of the alternating direction method of multipliers (ADMM). We empirically validate both variants of the algorithm, that we refer to as learned-SALSA (LSALSA), on image vision tasks and demonstrate that at inference our networks achieve vast improvements in terms of the running time and the quality of estimated sparse codes on both classic SC and MCA problems over more common baselines. We also demonstrate the visual advantage of our technique on the task of source separation. Finally, we present a theoretical framework for analyzing LSALSA network: we show that the proposed approach exactly implements a truncated ADMM applied to a new, learned cost function with curvature modified by one of the learned parameterized matrices. We extend a very recent stochastic alternating optimization analysis framework to show that a gradient descent step along this learned loss landscape is equivalent to a modified gradient descent step along the original loss landscape. In this framework, the acceleration achieved by LSALSA could potentially be explained by the network's ability to learn a correction to the gradient direction of steeper descent.

Keywords Sparse coding · Morphological component analysis · Deep learning

Editors: Karsten Borgwardt, Po-Ling Loh, Evimaria Terzi, Antti Ukkonen.

Electronic supplementary material The online version of this article (<https://doi.org/10.1007/s10994-019-05812-3>) contains supplementary material, which is available to authorized users.

✉ Benjamin Cowen
ben.cowen@nyu.edu

¹ New York University Tandon School of Engineering, 6 MetroTech Center, Brooklyn, NY 11201, USA

1 Introduction

In the SC framework, we seek to efficiently represent data by using only a sparse combination of available basis vectors. We therefore assume that an M -dimensional data vector $\mathbf{y} \in \mathbb{R}^M$ can be approximated as

$$\mathbf{y} \approx \mathbf{A}\mathbf{x}^*, \quad (1)$$

where $\mathbf{x}^* \in \mathbb{R}^N$ is sparse and $\mathbf{A} \in \mathbb{R}^{M \times N}$ is a dictionary, sometimes referred to as the synthesis matrix, whose columns are the basis vectors. This paper focuses on the generalized SC problem of decomposing a signal into morphologically distinct components. A typical assumption for this problem is that the data is a linear combination of D source signals:

$$\mathbf{y} = \sum_{i=1}^D \mathbf{y}_i. \quad (2)$$

The MCA framework (Starck et al. 2004) for addressing additive mixtures requires that each component \mathbf{y}_i admits a sparse representation within the corresponding dictionary \mathbf{A}_i , leading to a generalized signal approximation model:

$$\mathbf{y} \approx \sum_{i=1}^D \mathbf{A}_i \mathbf{x}_i^*. \quad (3)$$

We then seek to recover \mathbf{x}_i^* s given \mathbf{y} and dictionaries \mathbf{A}_i s. We may trivially satisfy (3) by setting, for example, $\mathbf{x}_i^* = 0$ for all $i \neq j$, and performing traditional SC using only dictionary \mathbf{A}_j . Thus, MCA further assumes that the dictionaries \mathbf{A}_i s are *distinct* in the sense that each source-specific dictionary allows obtaining sparse representation of the corresponding source signal, while being highly inefficient in representing the other content in the mixture. This assumption is difficult to enforce on harder problems, i.e. when the components \mathbf{y}_i have similar characteristics and do not admit intuitive a priori sparsifying bases. In practice, the \mathbf{A}_i s often have significant overlap in sparse representation, making the problem of jointly recovering the \mathbf{x}_i s highly ill-conditioned.

There exist iterative optimization algorithms for performing SC and MCA. The bottleneck of these techniques is that at inference a sparse code has to be computed for each data point or data patch (as in case of high-resolution images). In the single dictionary setting, ISTA (Daubechies et al. 2004) and FISTA (Beck and Teboulle 2009) are classical algorithmic choices for this purpose. For the MCA problem, the standard choice is SALSALSA (Afonso et al. 2011), an instance of ADMM (Boyd et al. 2011). The iterative optimization process is prohibitively slow for high-throughput real-time applications, especially in the case of the ill-conditioned MCA setting. Thus our goal is to provide algorithms performing efficient inference, i.e. algorithms that find good approximations of the optimal codes in significantly shorter time than FISTA or SALSALSA.

The first key contribution of this paper is an efficient and accurate deep learning architecture that is general enough to well-approximate optimal codes for both classic SC in a single-dictionary framework and MCA-based signal separation. By accelerating SALSALSA via learning, we provide a means for fast approximate source separation. We call our deep learning approximator Learned SALSALSA (LSALSALSA). The proposed encoder is formulated as a time-unfolded version of the SALSALSA algorithm with a fixed number of iterations, where the depth of the deep learning model corresponds to the number of SALSALSA iterations. We train the deep model in the supervised fashion to predict optimal sparse codes for a given input

and show that shallow architectures of fixed-depth, that correspond to only few iterations of the original SALSA, achieve superior performance to the classic algorithm.

The SALSA algorithm uses second-order information about the cost function, which gives it an advantage over popular comparators such as ISTA on ill-conditioned problems (Figueiredo et al. 2009). Our second key contribution is an empirical demonstration that this advantage carries over to the deep-learning accelerated versions LSALSA and LISTA (Gregor and LeCun 2010), while preserving SALSA's applicability to a broader class of learning problems such as MCA-based source separation (LISTA is used only in the single dictionary setting). To the best of our knowledge, our approach is the first one to utilize an instance of ADMM unrolled into a deep learning architecture to address a source separation problem

Our third key contribution is a theoretical framework that provides insight into how LSALSA is able to surpass SALSA, namely describing how the learning procedure can enhance the second-order information that is characteristically exploited by SALSA. In particular, we show that the forward-propagation of a signal through the LSALSA network is equivalent to the application of truncated-ADMM to a new, learned cost function, and present a theoretical framework for characterizing this function in relation to the original Augmented Lagrangian. To the best of our knowledge, our work is the first to attempt to analyze a learning-accelerated ADMM algorithm.

To summarize, our contributions are threefold:

1. We achieve significant acceleration in both SC and MCA: classic SALSA takes up to $100\times$ longer to achieve LSALSA's performance. This opens up the MCA framework to potentially be used in high-throughput, real-time applications.
2. We carefully compare an ADMM-based algorithm (SALSA) with our proposed learnable counterpart (LSALSA) and with popular baselines (ISTA and FISTA). For a large variety of computational constraints (i.e. fixed number of iterations), we perform comprehensive hyperparameter testing for each encoding method to ensure a fair comparison.
3. We present a theoretical framework for analyzing the LSALSA network, giving insight as to how it uses information learned from data to accelerate SALSA.

This paper is organized as follows: Sect. 2 provides literature review, Sect. 3 formulates the SC problem in detail, Sect. 4 shows how to derive predictive single dictionary SC and multiple dictionary MCA from their iterative counterparts and explains our approach (LSALSA). Section 5 elaborates our theoretical framework for analyzing LSALSA and provides insight into its empirically demonstrated advantages. Section 6 shows experimental results for both the single dictionary setting and MCA. Finally, Sect. 7 concludes the paper. We provide an open-source implementation of the sparse coding and source separation experiments presented herein.

2 Related work

A sparse code inference aims at computing sparse codes for given data and is most widely addressed via iterative schemes such as aforementioned ISTA and FISTA. Predicting approximations of optimal codes can be done using deep feed-forward learning architectures based on truncated convex solvers. This family of approaches lies at the core of this paper. A notable approach in this family known as LISTA (Gregor and LeCun 2010) stems from earlier predictive sparse decomposition methods (Kavukcuoglu et al. 2010; Jarrett et al. 2009), which however were obtaining approximations to the sparse codes of insufficient quality. LISTA

improves over these techniques and enhances ISTA by unfolding a fixed number of iterations to define a fixed-depth deep neural network that is trained with examples of input vectors paired with their corresponding optimal sparse codes obtained by conventional methods like ISTA or FISTA. LISTA was shown to provide high-quality approximations of optimal sparse codes with a fixed computational cost. Unrolling methodology has since been applied to algorithms solving SC with ℓ_0 -regularization (Wang et al. 2016) and message passing schemes (Borgerding and Schniter 2016). In other prior works, ISTA was recast as a recurrent neural network unit giving rise to a variant of LSTM (Gers et al. 2003; Zhou et al. 2018). Recently, theoretical analysis has been provided for LISTA (Chen et al. 2018; Moreau and Bruna 2016), in which the authors provide convergence analyses by imposing constraints on the LISTA algorithm. This analysis does not apply to the MCA problem as it cannot handle multiple dictionaries. In other words, they would approach the MCA problem by casting it as a SC problem with access to a single dictionary that is a concatenation of source-specific dictionaries, e.g. $[\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_D]$. Furthermore these analyses do not address the saddle-point setting as required for ADMM-type methods such as SALSAs.

MCA has been used successfully in a number of applications that include decomposing images into textures and cartoons for denoising and inpainting (Elad et al. 2005; Peyré et al. 2007, 2010; Shoham and Elad 2008; Starck et al. 2005a, b), detecting text in natural scene images (Liu et al. 2017), as well as other source separation problems such as separating non-stationary clutter from weather radar signals (Uysal et al. 2016), transients from sustained rhythmic components in EEG signals (Parekh et al. 2014), and stationary from dynamic components of MRI videos (Otazo et al. 2015). The MCA problem is frequently solved via SALSAs algorithm, which constitutes a special case of the ADMM method.

There exist a few approaches in the literature utilizing highly specialized trainable ADMM algorithms. One such framework (Yang et al. 2016) was demonstrated to improve the reconstruction accuracy and inference speed over a variety of state-of-the-art solvers for the problem of compressive sensing Magnetic Resonance Imaging. A variety of papers followed up on this work for various image reconstruction tasks, such as the Learned Primal-dual Algorithm (Adler and Öktem 2017). However, these approaches do not give a detailed iteration-by-iteration comparison of the baseline method versus the learned method, making it difficult to understand the accuracy/speed tradeoff. Another related framework (Sprechmann et al. 2013) was applied to efficiently learn task-specific (reconstruction or classification) sparse models via sparsity-promoting convolutional operators. None of the above methods were applied to the MCA or other source separation problems and moreover it is non-trivial to obtain such extensions of these works. An unrolled nonnegative matrix factorization (NMF) algorithm (Roux et al. 2015) was implemented as a deep network for the task of speech separation. In another work (Wisdom et al. 2017), the NMF-based speech separation task was solved with an ISTA-like unfolded network.

3 Problem formulation

This paper focuses on the inference problem in SC: given data vector \mathbf{y} and dictionary matrix \mathbf{A} , we consider algorithms for finding the unique coefficient vector \mathbf{x}^* that minimizes the ℓ_1 -regularized linear least squares cost function:

$$\mathbf{x}^* = \arg \min_{\mathbf{x}} \left\{ E_{\mathbf{A}}(\mathbf{x}; \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \alpha \|\mathbf{x}\|_1 \right\}, \quad (4)$$

where the scalar constant $\alpha \geq 0$ balances sparsity with data fidelity. Since this problem is convex, \mathbf{x}^* is unique and we refer to it as the optimal code for \mathbf{y} with respect to \mathbf{A} . The dictionary matrix \mathbf{A} is usually learned by minimizing a loss function given below (Olshausen and Field 1996)

$$\mathcal{L}_{\text{Dict}}(\mathbf{A}) = \frac{1}{P} \sum_{p=1}^P E_{\mathbf{A}}(\mathbf{x}^{*,p}; \mathbf{y}^p) \tag{5}$$

with respect to \mathbf{A} using stochastic gradient descent (SGD), where P is the size of the training data set, \mathbf{y}^p is the p th training sample, and $\mathbf{x}^{*,p}$ is the corresponding optimal sparse code. The optimal sparse codes in each iteration are obtained in this paper with FISTA. When training dictionaries, we require the columns of \mathbf{A} to have unit norm, as is common practice for regularizing the dictionary learning process (Olshausen and Field 1996), however this is not necessary for code inference.

In the MCA framework, a generalization of the cost function from Eq. 4 is minimized to estimate $\mathbf{x}_1^*, \mathbf{x}_2^*, \dots, \mathbf{x}_D^*$ from the model given in Eq. 3. Thus one minimizes

$$E_{\mathbf{A}}(\mathbf{x}; \mathbf{y}) = \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \sum_{i=1}^D \alpha_i \|\mathbf{x}_i\|_1, \tag{6}$$

using $\mathbf{A} := [\mathbf{A}_1, \mathbf{A}_2, \dots, \mathbf{A}_D] \in \mathbb{R}^{M \times N}$ and

$$\mathbf{x} := \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_D \end{bmatrix} \in \mathbb{R}^N, \tag{7}$$

where $\mathbf{x}_i \in \mathbb{R}^{N_i}$ for $i = \{1, 2, \dots, D\}$, $N = \sum_{i=1}^D N_i$, and α_i s are the coefficients controlling the sparsity penalties. We denote the concatenated optimal codes with $\mathbf{x}^* = \arg \min_{\mathbf{x}} E_{\mathbf{A}}(\mathbf{x}, \mathbf{y})$. To recover the single dictionary case, simply set $\alpha_i = \alpha_j, \forall i, j = 1, \dots, D$ and set \mathbf{A}_i to be partitions of \mathbf{A} .

In the classic MCA works, the dictionaries \mathbf{A}_i s are selected to be well-known filter banks with explicitly designed sparsification properties. Such hand-designed transforms have good generalization abilities and help to prevent overfitting. Also, MCA algorithms often require solving large systems of equations involving $\mathbf{A}^T \mathbf{A}$ or $\mathbf{A} \mathbf{A}^T$. An appropriate constraining of \mathbf{A}_i leads to a banded system of equations and in consequence reduces the computational complexity of these algorithms, e.g. Parekh et al. (2014). More recent MCA works use learned dictionaries for image analysis (Shoham and Elad 2008; Peyré et al. 2007). Some extensions of MCA consider learning dictionaries \mathbf{A}_i s and sparse codes jointly (Peyré et al. 2007, 2010).

Remark 1 (Learning dictionaries) In our paper, we learn dictionaries \mathbf{A}_i s independently. In particular, for each i we minimize

$$\mathcal{L}_{\text{Dict}}(\mathbf{A}_i) = \frac{1}{P} \sum_{p=1}^P E_{\mathbf{A}_i}(\mathbf{x}_i^{*,p}; \mathbf{y}_i^p) \tag{8}$$

with respect to \mathbf{A}_i using SGD, where \mathbf{y}_i^p is the i th mixture component of the p th training sample and $\mathbf{x}_i^{*,p}$ is the corresponding optimal sparse code. The columns are constrained to have unit norm. The sparse codes in each iteration are obtained with FISTA.

Algorithm 1 SALSA (Single Dictionary Afonso et al. 2010)

- 1: **Input:** $\alpha \geq 0, \mu > 0$
 $\mathbf{y} \in \mathbb{R}^M, \mathbf{A} \in \mathbb{R}^{M \times N}$
- 2: **Initialize:** $\mathbf{x} = \mathbf{A}^T \mathbf{y}$ and $\mathbf{d} = \mathbf{0}$
- 3: **repeat**
- 4: $\mathbf{u} = \text{soft}(\mathbf{x} + \mathbf{d}; \alpha/\mu)$
- 5: Solve for \mathbf{x} : $[\mu \mathbf{I} + \mathbf{A}^T \mathbf{A}] \mathbf{x} = \mathbf{A}^T \mathbf{y} + \mu(\mathbf{u} - \mathbf{d})$
- 6: $\mathbf{d} = \mathbf{d} - \mathbf{u} + \mathbf{x}$
- 7: **until** change in \mathbf{x} below a threshold

4 From iterative to predictive SC and MCA

4.1 Split augmented lagrangian shrinkage algorithm (SALSA)

The objective functions used in SC (Eq. 4) and MCA (Eq. 6) are each convex with respect to \mathbf{x} , allowing a wide variety of optimization algorithms with well-studied convergence results to be applied (Bauschke and Combettes 2011). Here we describe a popular algorithm that is general enough to solve both problems called SALSA (Afonso et al. 2010), which is an instance of ADMM. ADMM (Boyd et al. 2011) addresses an optimization problem with the form

$$\min_{\mathbf{x}} f_1(\mathbf{x}) + f_2(\mathbf{x}) \tag{9}$$

by re-casting it as the equivalent, constrained problem

$$\min_{\mathbf{u}, \mathbf{x}} f_1(\mathbf{x}) + f_2(\mathbf{u}) \quad \text{such that } \mathbf{x} = \mathbf{u}. \tag{10}$$

ADMM then optimizes the corresponding scaled Augmented Lagrangian,

$$\mathcal{L}_A = f_1(\mathbf{x}) + f_2(\mathbf{u}) + \frac{\mu}{2} \|\mathbf{u} - \mathbf{x} - \mathbf{d}\|_2^2 - \frac{\mu}{2} \|\mathbf{d}\|_2^2, \tag{11}$$

where \mathbf{d} correspond to Lagrangian multipliers, one variable at a time until convergence.

SALSA, proposed in Afonso et al. (2010), addresses an instance of the general optimization problem from Eq. 10 for which convergence has been proved in Eckstein and Bertsekas (1992). Namely, SALSA requires that (1) f_1 is a least-squares term, and (2) the proximity operator of f_2 can be computed exactly. For our most general cost function in Eq. 6, requirement (1) is clearly satisfied, and our f_2 is the weighted sum of ℓ_1 norms. In Supplemental Section A, we show that the the proximity operator of f_2 reduces to element-wise soft thresholding for each component, which in scalar form is given by

$$\text{soft}(z; \alpha) = \begin{cases} z - \alpha, & z > \alpha \\ 0, & |z| \leq \alpha \\ z + \alpha, & z < -\alpha \end{cases}. \tag{12}$$

When applied to a vector, $\text{soft}(\mathbf{z}; \alpha)$ performs soft thresholding element-wise. Thus, SALSA is guaranteed to converge for the multiple-dictionary sparse coding problem.

SALSA is given in Algorithms 1 and 2 for the single-dictionary case and the MCA case involving two dictionaries,¹ respectively. Note that in Algorithm 2, the \mathbf{u} and \mathbf{d} updates can

¹ In this paper we consider the MCA framework with two dictionaries. Extensions to more than two dictionaries are straightforward.

Algorithm 2 SALSA (Two Dictionaries Selesnick 2014)

- 1: **Input:** $\alpha_1, \alpha_2 \geq 0, \mu > 0$
 $y \in \mathbb{R}^M, \mathbf{A} \in \mathbb{R}^{M \times (N_1 + N_2)}$
- 2: **Initialize:** $\mathbf{x} = \mathbf{A}^T \mathbf{y}, \mathbf{d} := \begin{bmatrix} \mathbf{d}_1 \\ \mathbf{d}_2 \end{bmatrix} = 0$
- 3: **repeat**
- 4: $\mathbf{u} = \begin{bmatrix} \text{soft}(\mathbf{x}_1 + \mathbf{d}_1; \alpha_1/\mu) \\ \text{soft}(\mathbf{x}_2 + \mathbf{d}_2; \alpha_2/\mu) \end{bmatrix}$
- 5: Solve for \mathbf{x} :
 $\begin{bmatrix} \mu \mathbf{I} + \mathbf{A}^T \mathbf{A} \end{bmatrix} \mathbf{x} = \mathbf{A}^T \mathbf{y} + \mu(\mathbf{u} - \mathbf{d})$
- 6: $\mathbf{d} = \mathbf{d} - \mathbf{u} + \mathbf{x}$
- 7: **until** change in \mathbf{x} below a threshold




Fig. 1 A block diagram of SALSA. The one-time initialization $\mathbf{x} = \mathbf{A}^T \mathbf{y}$ is represented by a gate on the left

be performed with element-wise operations. The \mathbf{x} -update, however, is non-separable with respect to components $\{\mathbf{x}_i\}_{i=1}^D$ for general \mathbf{A} ; the system of equations in the \mathbf{x} -update cannot be broken down into D sub-problems, one for each component (in contrast, 1st order methods such as FISTA update components independently). We call this the *splitting step*.

As mentioned in Sect. 3, the \mathbf{x} -update is often simplified to element-wise operations by constraining matrix \mathbf{A} to have special properties. For example: requiring $\mathbf{A}\mathbf{A}^T = \rho \mathbf{I}, \rho \in \mathbb{R}_+$, reduces the \mathbf{x} -update step to element-wise division (after applying the matrix inverse lemma). In Yang et al. (2016), \mathbf{A} is set to be the partial Fourier transform, reducing the system of equations of the \mathbf{x} -update to be a series of convolutions and element-wise operations. In our work, as is typical in the case of SC, \mathbf{A} is a learned dictionary without any imposed structure.

Note that one way to solve for \mathbf{x} in Algorithms 1 and 2 is to compute the inverse of regularized Hessian matrix $\mu \mathbf{I} + \mathbf{A}^T \mathbf{A}$. This however needs to be done just once, at the very beginning, as this matrix remains fixed during the entire run of SALSA. We abbreviate the inverted matrix as

$$\mathbf{S} = (\mu \mathbf{I} + \mathbf{A}^T \mathbf{A})^{-1}. \tag{13}$$

We call this matrix a *splitting operator*. Note that the inversion process couples together the dictionary elements (and hence also the dictionaries) in a non-linear fashion. This is an advanced utilization of prior knowledge not seen in the comparator methods of Sect. 6. The recursive block diagram of SALSA is depicted in Fig. 1.

4.2 Learned SALSA (LSALSA)

We now describe our proposed deep encoder architecture that we refer to as Learned SALSA (LSALSA). Consider truncating the SALSA algorithm to a fixed number of iterations




Fig. 2 The deep learning architecture of LSALSA for $T = 3$. The soft-thresholding function, defined in Eq. 12, is an activation function found in each layer of the network and at the end

Algorithm 3 LSALSA forward pass (single dictionary)

```

Input:  $\alpha \geq 0, \mu > 0, \mathbf{y} \in \mathbb{R}^M$ 
 $\mathbf{x}(0) = \mathbf{W}_e \mathbf{y}, \mathbf{d}(0) = \mathbf{0}$ 

for  $t = 1$  to  $T$  do
     $\mathbf{u}(t) = \text{soft}(\mathbf{x}(t-1) + \mathbf{d}(t-1); \alpha/\mu)$ 
     $\mathbf{x}(t) = \mathbf{S}(\mathbf{W}_e \mathbf{y} + \mu(\mathbf{u}(t) - \mathbf{d}(t-1)))$ 
     $\mathbf{d}(t) = \mathbf{d}(t-1) - \mathbf{u}(t) + \mathbf{x}(t)$ 
end for

Output:  $\text{soft}(\mathbf{x}(t); \alpha/\mu)$ 
    
```

T and then time-unfolding it into a deep neural network architecture that matches the truncated SALSA’s output exactly. The obtained architecture is illustrated in Fig. 2 for $T = 3$, and the formulas for the t th layer w.r.t. the $(t - 1)$ th iterates are described via pseudocode in Algorithms 3 and 4 for the single-dictionary and MCA cases, respectively. Note that Algorithms 2 and 4 are the most general algorithms considered by us whereas Algorithms 1 and 3 are their special, i.e. single-dictionary, cases.

The LSALSA model has two matrices of learnable parameters: \mathbf{S} and \mathbf{W}_e . We initialize these to achieve an exact correspondence with SALSA:

$$\mathbf{W}_e = \mathbf{A}^T \in \mathbb{R}^{N \times M} \quad \text{and} \quad \mathbf{S} = (\mu \mathbf{I} + \mathbf{A}^T \mathbf{A})^{-1} \in \mathbb{R}^{N \times N}, \tag{14}$$

where $N = N_1 + N_2$ in the MCA case. All splitting operators \mathbf{S} share parameters across the network. LSALSA’s two matrices of parameters can be trained with standard backpropagation. Let $\mathbf{x} = f_e(\mathbf{W}_e, \mathbf{S}, \mathbf{y})$ denote the output of the LSALSA architecture after a forward propagation of \mathbf{y} . The cost function used for training the model is defined as

$$\mathcal{L}(\mathbf{W}_e, \mathbf{S}) = \frac{1}{2P} \sum_{p=1}^P \|\mathbf{x}^{*,p} - f_e(\mathbf{W}_e, \mathbf{S}, \mathbf{y}^p)\|_2^2. \tag{15}$$

To summarize, LSALSA extends SALSA. SALSA is meant to run until convergence, where LSALSA is meant to run for T iterations, where T is the depth of the network. Intuitively, the backpropagation steps applied during training in LSALSA fine-tune the “splitting step” so that T iterations can be sufficient to achieve good-quality sparse codes (those are obtained due to the existence of nonlinearities). The SALSA algorithm relies on cumulative Lagrange Multiplier updates to “explain away” code components, while separating sources. This is especially important in MCA, where similar atoms from different dictionaries will compete to represent the same segment of a mixed signal. The Lagrange Multiplier updates translate to a cross-layer connectivity pattern in the corresponding LSALSA network (see the d -updates in Fig. 2), which has been shown to be a beneficial architectural feature in e.g.

Algorithm 4 LSALSA forward pass (two dictionaries)

Input: $\alpha_1, \alpha_2 \geq 0, \mu > 0, \mathbf{y} \in \mathbb{R}^M$
 $\mathbf{x}(0) = \mathbf{W}_e \mathbf{y}, \mathbf{d}(0) = \mathbf{0}$

for $t = 1$ **to** T **do**
 $\mathbf{u}(t) = \begin{bmatrix} \text{soft}(\mathbf{x}_1(t-1) + \mathbf{d}_1(t-1); \alpha_1/\mu) \\ \text{soft}(\mathbf{x}_2(t-1) + \mathbf{d}_2(t-1); \alpha_2/\mu) \end{bmatrix}$
 $\mathbf{x}(t) = \mathbf{S}(\mathbf{W}_e \mathbf{y} + \mu(\mathbf{u}(t) - \mathbf{d}(t-1)))$
 $\mathbf{d}(t) = \mathbf{d}(t-1) - \mathbf{u}(t) + \mathbf{x}(t)$
end for

Output: $\begin{bmatrix} \text{soft}(\mathbf{x}_1(t); \alpha_1/\mu) \\ \text{soft}(\mathbf{x}_2(t); \alpha_2/\mu) \end{bmatrix}$

(Greff et al. 2016; Liao and Poggio 2016; Orhan and Pitkow 2018). During training, LSALSA is fine-tuning the splitting operator \mathbf{S} so that it need not rely on a large number of cumulative updates. However, we show in Sect. 5 that even after training, forward propagation through an LSALSA network is equivalent to the application of a truncated ADMM algorithm applied to a new, learned cost function that generalizes the original problem.

5 Analysis of LSALSA

5.1 Optimality property for LSALSA

Typically, analyses of ADMM-like algorithms rely on the optimality of each primal update, e.g. that $\mathbf{x}^{(k+1)} = \arg \min_{\mathbf{x}} \mathcal{L}_A(\mathbf{x}, \mathbf{u}^{(k+1)}; \mathbf{d}^{(k)})$ (Boyd et al. 2011; Goldstein et al. 2014; Wang et al. 2019). In Theorem 1 we show that LSALSA provides optimal primal updates with respect to a generalization of the Augmented Lagrangian (11) parameterized by \mathbf{S} . The proof is provided in Supplemental Section C.

Theorem 1 (LSALSA Optimality) *Given a neural network with the LSALSA architecture as described in Sect. 4.2, there exists an Augmented Lagrangian for which the LSALSA network provides optimal primal updates. In particular, for learned matrices \mathbf{S} and \mathbf{W}_e , we have*

$$\hat{\mathcal{L}}_A = \hat{f}_1(\mathbf{x}; \mathbf{S}) + \ell_1(\mathbf{u}) + \frac{\mu}{2} \|\mathbf{u} - \mathbf{x} - \mathbf{d}\|^2 - \frac{\mu}{2} \|\mathbf{d}\|^2, \tag{16}$$

where

$$\hat{f}_1(\mathbf{x}; \mathbf{S}) = \frac{1}{2} \mathbf{x}^T [\mathbf{S}^{-1} - \mu \mathbf{I}] \mathbf{x} - (\mathbf{W}_e \mathbf{y})^T \mathbf{x} + \frac{1}{2} \mathbf{y}^T \mathbf{y}, \tag{17}$$

and $\ell_1(\mathbf{u})$ represents a sum of L1-terms as in (6).

Remark 2 (LSALSA as an instance of ADMM) Note that by plugging in the initializations of \mathbf{S} and \mathbf{W}_e , given in Eq. 14, we recover the original Augmented Lagrangian. Then, from the perspective of Theorem 1, LSALSA at inference is equivalent to applying T iterations of ADMM on a new, learned cost function that generalizes the original problem in Eq. 11.

Remark 3 (LSALSA provides sparse solutions) Since $\hat{\mathcal{L}}_A$ employs the ℓ_1 -norm in the usual way and LSALSA’s \mathbf{u} -update is standard soft-thresholding, we can expect LSALSA to enforce sparsity given sufficient iterations.

We show in Sect. 5.2 that the optimal direction for $\hat{\mathcal{L}}_A$ is related to the optimal direction for \mathcal{L}_A , and in Sect. 5.3 we show that gradient descent along $\hat{\mathcal{L}}_A$ is equivalent to a modified gradient descent along \mathcal{L}_A . For simplicity, we consider the case of learned, symmetric \mathbf{S} while holding fixed $\mathbf{W}_e \equiv \mathbf{A}^T$.

5.2 Modified descent direction: deterministic framework

Though $\hat{\mathcal{L}}_A$'s dependence on \mathbf{u} and \mathbf{d} is standard in ADMM settings (Boyd et al. 2011), the learned data-fidelity term \hat{f}_1 that commands \mathbf{x} -directions is now a data-driven quadratic form that relies on the weight matrix \mathbf{S} that parameterizes LSALSA. We will next rewrite the new cost function in terms of the original Augmented Lagrangian:

$$\hat{\mathcal{L}}_A(\mathbf{x}, \mathbf{u}, \mathbf{d}) = \mathcal{L}_A(\mathbf{x}, \mathbf{u}, \mathbf{d}) + \hat{f}_1(\mathbf{x}; \mathbf{S}) - \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2. \tag{18}$$

The optimality condition for $\hat{\mathcal{L}}_A$ can be written

$$\begin{aligned} 0 &= \nabla_{\mathbf{x}} \hat{\mathcal{L}}_A(\mathbf{x}^*, \mathbf{u}, \mathbf{d}) \\ &= \nabla_{\mathbf{x}} \left(\mathcal{L}_A(\mathbf{x}^*, \mathbf{u}, \mathbf{d}) + \hat{f}_1(\mathbf{x}; \mathbf{S}) - \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \right) \\ &= \nabla_{\mathbf{x}} \mathcal{L}_A(\mathbf{x}^*, \mathbf{u}, \mathbf{d}) + [\mathbf{S}^{-1} - \mu I - \mathbf{A}^T \mathbf{A}] \mathbf{x}^*. \end{aligned}$$

Then, using $\nabla_{\mathbf{x}}^2 \mathcal{L}_A = \mu I + \mathbf{A}^T \mathbf{A}$ we can write the LSALSA update as

$$0 = \nabla_{\mathbf{x}} \mathcal{L}_A(\mathbf{x}^*, \mathbf{u}, \mathbf{d}) + [\mathbf{S}^{-1} - \nabla_{\mathbf{x}}^2 \mathcal{L}_A] \mathbf{x}^* \tag{19}$$

$$\Rightarrow [\mathbf{S}^{-1} - \nabla_{\mathbf{x}}^2 \mathcal{L}_A] \mathbf{x}^* = -\nabla_{\mathbf{x}} \mathcal{L}_A(\mathbf{x}^*, \mathbf{u}, \mathbf{d}). \tag{20}$$

The root-finding problem posed in (19) and equivalent system of equations in (20) resemble a Newton-like update, but using a learned modification of the original Lagrangian's Hessian matrix. Note that at initialization (using Formula 14), the left-hand-side cancels to zero, recovering the optimality condition for the original problem. This also admits an intuition that LSALSA is incorporating prior knowledge, learned from the training data, that could be made to balance between optimality of the original problem while maintaining some relationship with the training data distribution.

5.3 Modified descent direction: stochastic framework

We will next look at (L)SALSA through the prism of worst-case analysis, i.e. by replacing the optimal primal steps with stochastic gradient descent. This effectively enables us to analyze (L)SALSA as a stochastic alternated optimization approach solving a general saddle point problem, and we show that LSALSA leads to faster convergence under certain assumptions that we stipulate. Our analysis is a direct extension of that in Choromanska et al. (2019). We provide the final statement of the theorem below and defer all proofs to the supplement.

5.3.1 Problem formulation

Consider the following general saddle-point problem:

$$\max_{\phi_1, \dots, \phi_{K_2}} \min_{\theta_1, \dots, \theta_{K_1}} \mathcal{L}(\theta_1, \dots, \theta_{K_1}; \phi_1, \dots, \phi_{K_2}) \tag{21}$$

$$\begin{aligned} & \Updownarrow \\ & \max_{\phi} \min_{\theta} \mathcal{L}(\theta; \phi), \end{aligned} \tag{22}$$

using $\theta = [\theta_1, \dots, \theta_{K_1}]$ to denote the collection of variables to be minimized, and $\phi = [\phi_1, \dots, \phi_{K_2}]$ the variables to be maximized. We denote the entire collection of variables as $\mathbf{x} = [\theta, \phi] \in \mathbb{R}^K$, where $K = K_1 + K_2$ is the total number of arguments. We denote with x_d the d th entry in \mathbf{x} . For theoretical analysis we consider a smooth function \mathcal{L} as is often done in the literature (especially for ℓ_1 problems, as discussed in Lange et al. 2014; Schmidt et al. 2007).

Let (x_1^*, \dots, x_K^*) be the optimal solution of the saddle point problem in (22), where \mathcal{L} is computed over global data population (i.e. averaged over an infinite number of samples). For each variable x_d , we assume a lower bound on the radii of convergence $r_d > 0$. Let $\nabla_d^1 \mathcal{L}$ denote the gradient of \mathcal{L} with respect to the d th argument evaluated on a single data sample (stochastic gradient), and $\nabla_d \mathcal{L}$ to be that with respect to the global data population (i.e. an “oracle gradient”).

We analyze an Alternating Optimization algorithm that, at the d th step, optimizes \mathcal{L} with respect to x_d while holding all other $x_{i \neq d}$ fixed:

$$x_d^{t+1} = \Pi_d(x_d^t \pm \eta^t \nabla_d^1 \mathcal{L}_{x_d}^t), \tag{23}$$

using the \pm symbol to denote gradient descent for $d \leq K_1$ and gradient ascent for $d > K_1$. Π_d is the projection onto the Euclidean-ball $B_2(\frac{r_d}{2}, x_d^*)$, with radius $\frac{r_d}{2}$ and centered around the optimal value x_d^* : this ensures that for each d , all iterates of x_d remain within the r_d -ball around x_d^* .²

5.3.2 Assumptions

The following assumptions are necessary for the Theorems in Sect. 5.3.3. The mathematical definitions of strong-convexity, strong-concavity, and smoothness follow the standards from Nesterov (2013).

Assumption 1 (Convex–Concave) For each $d \leq K_1$, $\mathcal{L}_{x_d}^*$ is β_d -convex, and for each $d > K_1$, $\mathcal{L}_{x_d}^*$ is β_d -concave within a ball around the solution x_d^* of radius r_d .

Assumption 2 (Smoothness) For all $d \in \{1, \dots, K\}$, the function $\mathcal{L}_{x_d}^*$ is α_d -smooth.

In summary, for every $d = 1, \dots, K$, $\mathcal{L}_{x_d}^*$ is either β_d -convex or concave in a neighborhood around the optimal point, and α_d -smooth. Next we assume two standard properties on the gradient of the cost function.

Assumption 3 (Gradient Stability GS(γ_d)) We assume that for each $d = 1, \dots, K$, the following gradient stability condition holds for $\gamma_d \geq 0$ over the Euclidean ball $x_d \in B_2(r_d, x_d^*)$:

$$\|\nabla_d \mathcal{L}_{x_d}^* - \nabla_d \mathcal{L}_{x_d}\| \leq \gamma_d \sum_{i \neq d} \|x_i - x_i^*\|. \tag{24}$$

Assumption 4 (Assumption A.6: Bounded Gradient) We assume that the expected value of the gradient of our objective function \mathcal{L} is bounded by $\sigma = \sqrt{\sum_{d=1}^K \sigma_d^2}$, where:

$$\sigma_d = \sup \left\{ \mathbb{E} \left[\|\nabla_d \mathcal{L}_{x_d}\|^2 \right] : x_d \in B_2(r_d, x_d^*), \forall d = 1, \dots, K \right\}. \tag{25}$$

² this assumption can be potentially eliminated with carefully selected initial stepsizes.

5.3.3 Convergence statement

Denote with $\Delta_d^t = x_d^t - x_d^*$ the error of the t th estimate of d th element of the global optimizer \mathbf{x}^* . Define the following:

$$\mathcal{E}_{\text{SALSA}}(\beta) = \left(\frac{2}{t+3}\right)^{\frac{3}{2}} \mathbb{E} \left[\sum_{d=1}^K \|\Delta_d^0\|^2 \right] + \frac{9\sigma^2}{[2\xi(\beta) - \gamma(2K-1)]^2(t+3)}, \quad (26)$$

where $\xi(\beta)$ increases monotonically with increasing β .

Theorem 2 (Convergence of SALSA and LSALSA) *Suppose that cost functions underlying SALSA \mathcal{L}_A and LSALSA $\hat{\mathcal{L}}_A$ satisfy the Assumptions in Sect. 5.3.2 with convexity moduli β and $\hat{\beta}$ (the latter is implicitly learned from the data). Assume also that the deep model representing LSALSA had enough capacity to learn $\hat{\beta}$ such that $\hat{\beta} > \beta$, while keeping the same location of the global optimal fixed point, \mathbf{x}^* .³*

Then, using the Stochastic Alternating Optimization scheme in Eq. 23 on \mathcal{L}_A and $\hat{\mathcal{L}}_A$ such that the requirements from Theorem 4 are satisfied, starting from the same initial point, the error satisfies the following:
for SALSA:

$$\sum_{d=1}^K \|\Delta_d^{t+1}\|^2 \leq \mathcal{E}_{\text{SALSA}}(\beta), \quad (27)$$

and for LSALSA:

$$\sum_{d=1}^K \|\Delta_d^{t+1}\|^2 \leq \mathcal{E}_{\text{LSALSA}}(\hat{\beta}) = \mathcal{E}_{\text{SALSA}}(\beta) - \Delta_\beta, \quad (28)$$

where

$$\Delta_\beta = \mathcal{O} \left(\frac{\hat{\beta}^2 - \beta^2}{(2\beta\hat{\beta})^2} \right). \quad (29)$$

The above theorem states that, given enough capacity of the deep model, LSALSA can learn steeper descent direction than SALSA. We provide below an intuition for that. Consider the gradient descent step (or its stochastic approximation) for $\hat{\mathcal{L}}_A$ in the \mathbf{x} -direction as given below

$$\begin{aligned} \mathbf{x}^{(k+1)} &= \mathbf{x}^{(k)} - \eta^k \nabla_{\mathbf{x}} \hat{\mathcal{L}}_A(\mathbf{x}^{(k)}, \mathbf{u}^{(k+1)}, \mathbf{d}^{(k)}) \\ &= \mathbf{x}^{(k)} - \eta^k \nabla_{\mathbf{x}} \left(\mathcal{L}_A^k + \phi(\mathbf{x}; \mathbf{S}) - \frac{1}{2} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 \right) \\ &= \mathbf{x}^{(k)} - \eta^k \nabla_{\mathbf{x}} \mathcal{L}_A^k - \eta^k [\mathbf{S}^{-1} - \mu I - \mathbf{A}^T \mathbf{A}] \mathbf{x}^{(k)} \\ &= \underbrace{\mathbf{x}^{(k)} - \eta^k \nabla_{\mathbf{x}} \mathcal{L}_A^k}_{\text{unlearned descent step}} - \eta^k [\mathbf{S}^{-1} - \nabla_{\mathbf{x}}^2 \mathcal{L}_A] \mathbf{x}^{(k)} \\ &= [I - \eta^k P] \mathbf{x}^{(k)} - \eta^k \nabla_{\mathbf{x}} \mathcal{L}_A^k, \end{aligned} \quad (30)$$

where $P := \mathbf{S}^{-1} - \nabla_{\mathbf{x}}^2 \mathcal{L}_A$.

³ LSALSA is trained to keep the same global fixed point, see Eq. 15.

This update can be seen as taking first a gradient descent step and then pushing the optimizer further in the learned direction, which we empirically show is a faster direction of decent.

6 Numerical experiments

We now present a variety of sparse coding inference tasks to evaluate our algorithm's speed, accuracy, and sparsity trade-offs. For each task (including both SC and MCA), we consider a variety of settings of T , i.e. the number of iterations, and do a full hyperparameter grid search for each setting. In other words, we ask "how well can each encoding algorithm approximate the optimal codes, given a fixed number of stages?". We compare LSALSA, truncated SALSA, truncated FISTA, and LISTA (Gregor and LeCun 2010) in terms of their RMSE proximity to optimal codes, sparsity levels, and performance on classification tasks. Both LSALSA and LISTA are implemented as feedforward neural networks. For MCA experiments, we run FISTA and LISTA using the concatenated dictionary \mathbf{A} .

We focus on the inference problem and thus learn the dictionaries off-line as described in Sect. 3. Dictionary learning is performed only once for each data set, and the resulting dictionaries are held constant across all methods and experiments herein (visualization of the atoms of the obtained dictionaries can be found in Section F in the Supplement). For MCA, the independently-learned dictionaries are still used, creating difficult ill-conditioned problems (because each dictionary is able to at least partially represent both components).

To train the encoders, we minimize Eq. 15 with respect to \mathbf{W}_e and \mathbf{S} using vanilla Stochastic Gradient Descent (SGD). We considered the optimization complete after a fixed number of epochs, or when the relative change in cost function fell below a threshold of 10^{-6} . During hyperparameter grid searches, only 10 epochs through the training data were allowed; for testing, 100 epochs of training were allowed (usually the tolerance was reached before 100 epochs). The optimal codes are determined prior to training by solving the convex inference problem with fixed α^* and μ^* , e.g. by running FISTA or SALSA to convergence (details are discussed in each section). In order to set the α^* , μ^* , we fix $\mu^* = 10$ and tune α^* to yield an average sparsity of at least 89%. We then slowly increase α^* s until just before the optimal sparse codes' fail to provide recognizable image reconstructions. We take the simplest approach to image reconstruction: simply multiplying the sparse code with its corresponding dictionary. No additional learning was performed to achieve reconstruction: i.e. for LSALSA we have $\mathbf{A}_i \cdot (f_e(\mathbf{W}_e, \mathbf{S}, \mathbf{y}))_i$, where $f_e(\mathbf{W}_e, \mathbf{S}, \mathbf{y})_i$ represents the i -th component of the encoder's output.

We implemented the experiments in Lua using Torch7, and executed the experiments on a 64-bit Linux machine with 32GB RAM, i7-6850K CPU at 3.6GHz, and GTX 1080 8GB GPU. The hyperparameters were selected via a grid search with specific values listed in the Supplement, Section E.

6.1 Single dictionary (SC) case

We run SC experiments with four data sets: Fashion MNIST (Xiao et al. 2017) (10 classes), ASIRRA (Elson et al. 2007) (2 classes), MNIST (LeCun et al. 2009) (10 classes), and CIFAR-10 (Krizhevsky and Hinton 2009) (10 classes). The ASIRRA data set is a collection of natural images of cats and dogs. We use a subset of the whole data set: 4000 training images and

1000 testing images as commonly done (Golle 2008). The results for MNIST and CIFAR-10 are reported in Section G in the Supplement.

The 32×32 Fashion MNIST images were first divided into 10×10 non-overlapping patches (ignoring extra pixels on two edges), resulting in 9 patches per image. Then, optimal codes were computed for each vectorized patch by minimizing the objective from Eq. 4 with FISTA for 200 iterations. The ASIRRA images come in varying sizes. We resized them to the resolution of 224×224 via Torch7's bilinear interpolation and converted each image to grayscale. Then we divided them into 16×16 non-overlapping patches, resulting in 196 patches per image. Optimal codes were computed patch-wise as for Fashion MNIST, but taking 700 iterations to ensure convergence on this more difficult SC problem. For Fashion MNIST we selected $\alpha^* = 0.15$ and for ASIRRA, $\alpha^* = 0.5$. using criteria mentioned earlier in the Section.

The data sets were then separated into training and testing sets. The training patches were used to produce the dictionaries. Visualizations of the dictionary atoms are provided in Section F in the Supplement. An exhaustive hyper-parameter search⁴ was performed for each encoding method and for each number of iterations T , to minimize RMSE between obtained and optimal codes. The hyper-parameter search included α for all methods, μ for SALSAs and LSALSAs, as well as SGD learning rates and learning rate decay schedules for LSALSAs and LISTAs training.

The obtained encoders were used to compute sparse codes on the test set. Those were then compared with the optimal codes via RMSE. The results for Fashion MNIST are shown both in terms of the number of iterations and the wallclock time in seconds used to make the prediction (Fig. 3). It takes FISTA more than 15 iterations and SALSAs more than 5 to reach the error achieved by LSALSAs in just one. Near $T = 100$, both FISTA and SALSAs are finally converging to the optimal codes. LISTA outperforms FISTA at first, but does not show much improvement after $T > 10$. Similar results for ASIRRA are shown in the same figure. On this more difficult problem, it takes FISTA more than 50 iterations and SALSAs more than 20 to catch up with LSALSAs with a single iteration. LISTA and LSALSAs are comparable for $T \leq 5$, after which LSALSAs dramatically improves its optimal code prediction and, similarly as in case of Fashion MNIST, shows advantage in terms of the number of iterations, inference time, and the quality of the recovered sparse codes over other methods.

We also investigated which method yields better codes in terms of a classification task. We trained a logistic regression classifier to predict the label from the corresponding optimal sparse code, then ask: "can the classifier still recognize a fast encoder's estimate to the optimal code?". For Fashion MNIST each image is associated with 9 optimal codes (one for each patch), yielding a total feature length of $9 \times 10 \times 10 = 900$. The Fashion MNIST classifier was trained until it achieved 0% classification error on the optimal codes. For ASIRRA, each concatenated optimal code had length $196 \times 16 \times 16 = 50,176$; to reduce the dimensionality we applied a random Gaussian projection $\mathcal{G} : \mathbb{R}^{50,176} \rightarrow \mathbb{R}^{500}$ before inputting the codes into the classifier. The classifier was trained on the optimal projected codes of length 500 until it achieved 0.5% error. The results for Fashion MNIST and ASIRRA are shown in Table 3 and 4, respectively, in Section G in the Supplement. *Note:* The classifier was trained on the target test codes so that the resulting classification error is only due to the difference between the optimal and estimated codes. In conclusion, although the FISTA, LISTA, or SALSAs codes may not look that much worse than LSALSAs in terms of RMSE, we see in the Tables that the expert classifiers cannot recognize the extracted codes, despite being trained to recognize the optimal codes which the algorithms seek to approximate.

⁴ The parameter settings that we explored in all our experiments are provided in the Supplement.

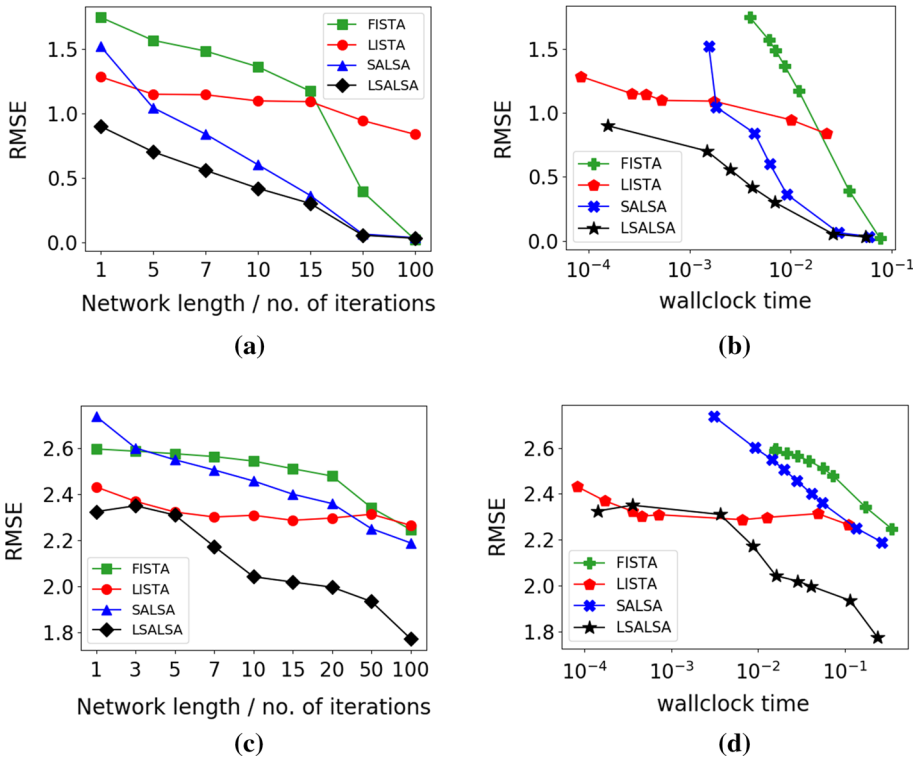


Fig. 3 Code prediction error as a function **a** iteration count, and **b** inference wallclock time for Fashion MNIST (a, b) and ASIRRA (c, d)

6.2 MCA: two-dictionary case

6.2.1 Data preparation

We now describe the dataset that we curated for the MCA experiments. We address the problem of decoupling numerals (text) from natural images, a topic closely related to text detection in natural scenes (Liu et al. 2017; Tian et al. 2015; Yan et al. 2018). Following the notation introduced previously in the paper, we set y_1^p s to be the whole 32×32 MNIST images and y_2^p s to be non-overlapping 32×32 patches from ASIRRA (thus we have 49 patches per image). We obtain 196k training and 49k testing patches from ASIRRA, and 60k training and 10k testing images from MNIST. We add together randomly selected MNIST images and ASIRRA patches to generate 588k mixed training images and 49k mixed testing images. Optimal codes were computed using SALSA (Algorithm 2) for 100 iterations, ensuring that each component had a sparsity level greater than 89%, while retaining visually recognizable reconstructions. The values selected were $\alpha_1 = 0.125^*$, $\alpha_2^* = 0.2$, $\mu^* = 10$. We also performed MCA experiments on additive mixtures of CIFAR-10 and MNIST images. Those results can be found in Section H in the Supplement.

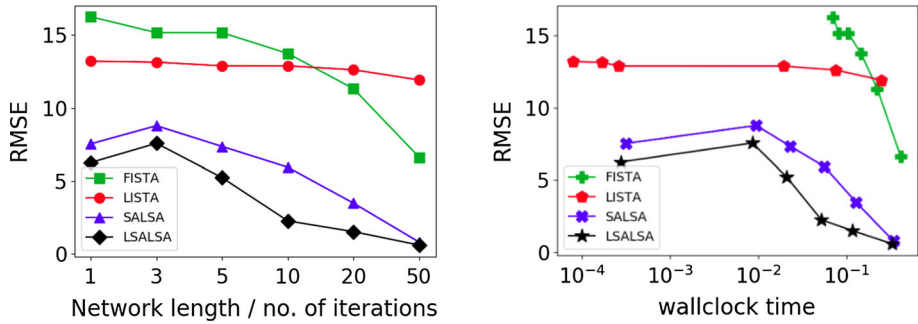


Fig. 4 MCA experiment using MNIST + ASIRRA data set. (left) Code prediction errors for varying numbers of iterations. (right) Code prediction error versus inference wallclock time

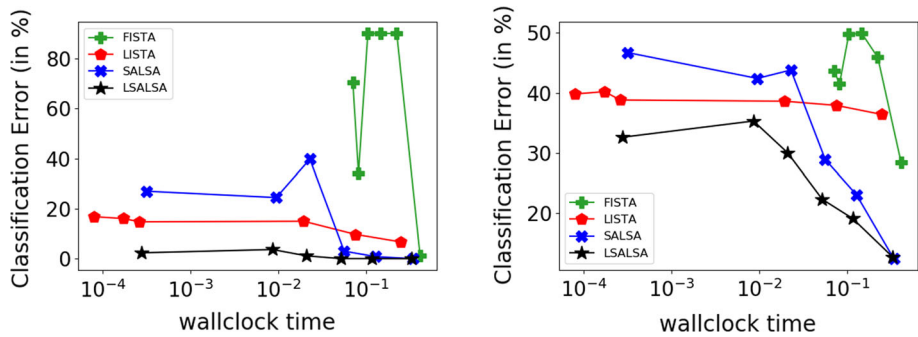


Fig. 5 MCA experiment separating MNIST + ASIRRA components: the trade-off between the sparse codes classification error versus their inference time is captured for different network lengths on (left) for MNIST (right) for ASIRRA

6.2.2 Results

An exhaustive hyper-parameter search was performed for each encoding method and each number of iterations T . The hyper-parameters search included α for FISTA and LISTA, α_1, α_2, μ for SALSA and LSALSA, as well as SGD learning rates for LSALSA and LISTA training. The code prediction error curves are presented in Fig. 4. LSALSA steadily outperforms the others, until SALSA catches up around $T = 50$. FISTA and LISTA, without a mechanism for distinguishing two dictionaries, struggle to estimate the optimal codes (Fig. 5).

In Fig. 6 we illustrate each method’s sparsity/accuracy trade-off on the ASIRRA test data set, while varying T (Supplemental Section I contains a similar plot for MNIST). For each data point in the test set, we plot its sparsity versus RMSE code-error, resulting in a point-cloud for each algorithm. For example, a sparsity value of 0.6 corresponds to 60% of the code elements being equal to zero. These point clouds represent the tradeoff between sparsity and fidelity to the original targets (e.g. proximity to the global solution as defined in original the convex problem). For each T , the (black) LSALSA point-cloud is generally further to the right and/or located below the other point-clouds, representing higher sparsity and/or lower error, respectively. For example, while FISTA achieves some mildly sparser solutions for $T = 10, 20$, it significantly sacrifices RMSE. In this sense, we argue that LSALSA enjoys the best sparsity-accuracy trade-off from among the four methods.

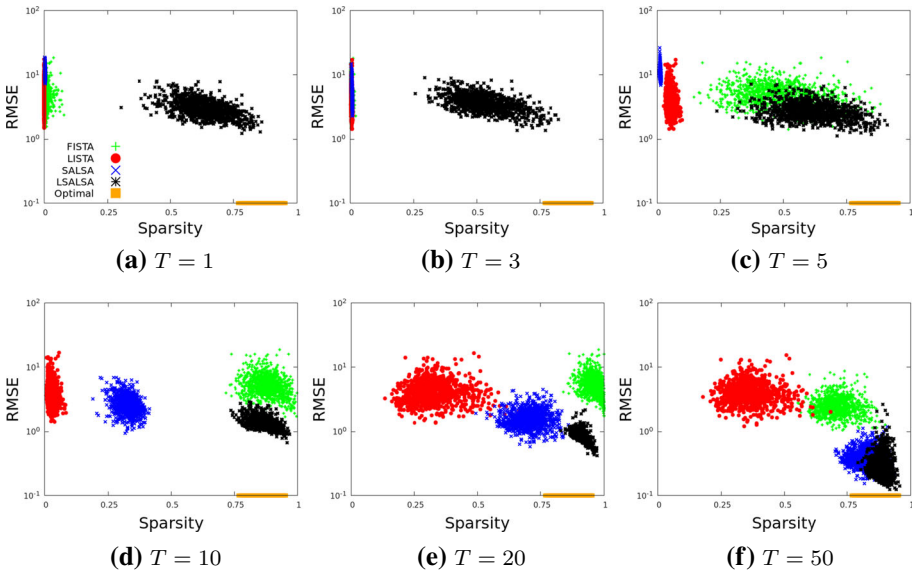


Fig. 6 Sparsity/accuracy trade-off analysis for ASIRRA obtained for the source separation experiment with MNIST + ASIRRA data set. Each method corresponds to a colored point cloud, where each point corresponds to one sample from the ASIRRA test data set. LSALSA (black) achieves the higher sparsity and/or lower code estimation error than the other methods for each T

Table 1 MNIST classification error obtained after source separation (10 classes). The best performer is in bold

Iter	Classification error (in %)			
	FISTA	LISTA	SALSA	LSALSA
1	70.29	16.81	27.00	2.37
3	34.09	16.10	24.45	3.69
5	89.97	14.78	39.74	1.15
10	90.00	15.00	3.03	0.05
20	90.00	9.74	0.85	0.05
50	1.30	6.73	0.02	0.02

Table 2 ASIRRA classification error obtained after source separation (2 classes). The best performer is in bold

Iter	Classification error (in %)			
	FISTA	LISTA	SALSA	LSALSA
1	43.70	39.80	46.70	32.60
3	41.50	40.20	42.40	35.30
5	49.80	38.80	43.80	30.00
10	49.90	38.60	28.90	22.30
20	45.50	37.90	23.00	19.10
50	28.40	36.40	12.40	12.70

Similarly as before, we performed an evaluation on the classification task. A separate classifier was trained for each data set using the separated optimal codes $\mathbf{x}_1^{*,P}$ and $\mathbf{x}_2^{*,P}$, respectively. As before, a random Gaussian projection was used to reduce the ASIRRA codes




Fig. 7 MCA experiment using MNIST + ASIRRA. Image reconstructions obtained by SALSA, LSALSA, FISTA, LISTA for $T = 1, 5$. Top row: original data (components and mixed)

to the length 500 before inputting to the classifier. The classification results are depicted in Table 1 for MNIST and Table 2 for ASIRRA.

Finally, in Fig. 7 we present exemplary reconstructed images obtained by different methods when performing source separation (more reconstruction results can be found in Section J in the Supplement). FISTA and LISTA are unable to separate components without severely

corrupting the ASIRRA component. LSALSA has visually recognizable separations even at $T = 1$, and the MNIST component is almost gone by $T = 5$. Recall that no additional learning is employed to generate reconstructions, they are simply codes multiplied by corresponding dictionary matrices.

7 Conclusions

In this paper we propose a deep encoder architecture LSALSA, obtained from time-unfolding the split augmented lagrangian shrinkage algorithm (SALSA). We empirically demonstrate that LSALSA inherits desired properties from SALSA and outperforms baseline methods such as SALSA, FISTA, and LISTA in terms of both the quality of predicted sparse codes, and the running time in both the single and multiple (MCA) dictionary case. In the two-dictionary MCA setting, we furthermore show that LSALSA obtains the separation of image components faster, and with better visual quality than the separation obtained by SALSA. The LSALSA network can tackle the general single and multiple dictionary coding problems without extension, unlike common competitors.

We also present a theoretical framework to analyze LSALSA. We show that the forward propagation of a signal through the LSALSA network is equivalent to a truncated ADMM algorithm applied to a new, learned cost function that generalizes the original problem. We show via the optimality conditions for this new cost function that the LSALSA update is related to a “learned pseudo-Newton” update down the original loss landscape, whose descent direction is corrected by a learned modification of the Hessian of the original cost function. Finally, we extend a very recent Stochastic Alternating Optimization analysis framework to show that a *gradient descent step down the learned loss landscape is equivalent with taking a modified gradient descent step along the original loss landscape*. In this framework we provide conditions under which LSALSA’s descent direction modification can speed up convergence.

References

- Adler, J., & Öktem, O. (2017). Learned primal-dual reconstruction. CoRR [arXiv:1707.06474](https://arxiv.org/abs/1707.06474).
- Afonso, M., Bioucas-Dias, J., & Figueiredo, M. (2010). Fast image recovery using variable splitting and constrained optimization. *IEEE Transactions on Image Processing*, 19(9), 2345–2356.
- Afonso, M., Bioucas-Dias, J., & Figueiredo, M. (2011). An augmented Lagrangian approach to the constrained optimization formulation of imaging inverse problems. *IEEE Transactions on Image Processing*, 20(3), 681–695.
- Bauschke, H. H., & Combettes, P. L. (2011). *Convex analysis and monotone operator theory in Hilbert spaces* (1st ed.). Berlin: Springer.
- Beck, A., & Teboulle, M. (2009). A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM: SIAM Journal on Imaging Sciences*, 2(1), 183–202.
- Borgerding, M., & Schniter, P. (2016) Onsager-corrected deep learning for sparse linear inverse problems. In *GlobalSIP*.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1), 1–122.
- Chen, X., Liu, J., Wang, Z., & Yin, W. (2018). Theoretical linear convergence of unfolded ISTA and its practical weights and thresholds. arXiv preprint [arXiv:1808.10038](https://arxiv.org/abs/1808.10038).
- Choromanska, A., Cowen, B., Kumaravel, S., Luss, R., Rish, I., Kingsbury, B., Tejwani, R., & Bouneffouf, D. (2019). Beyond backprop: Alternating minimization with co-activation memory. arXiv preprint [arXiv:1806.09077v3](https://arxiv.org/abs/1806.09077v3).

- Daubechies, I., Defrise, M., & De Mol, C. (2004). An iterative thresholding algorithm for linear inverse problems with a sparsity constraint. *Communications on Pure and Applied Mathematics*, 57(11), 1413–1457.
- Eckstein, J., & Bertsekas, D. (1992). On the Douglas–Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55, 293–318.
- Elad, M., Starck, J. L., Querre, P., & Donoho, D. L. (2005). Simultaneous cartoon and texture image inpainting using morphological component analysis (MCA). *Applied and Computational Harmonic Analysis*, 19(3), 340–358.
- Elson, J., Douceur, J., Howell, J., & Saul, J. (2007). Asirra: A CAPTCHA that exploits interest-aligned manual image categorization. In *ACM CCS*.
- Figueiredo, M., Bioucas-Dias, J., & Afonso, M. (2009). Fast frame-based image deconvolution using variable splitting and constrained optimization. In *Proceedings of IEEE workshop on statistical signal processing* (pp. 109–112).
- Gers, F. A., Schraudolph, N. N., & Schmidhuber, J. (2003). Learning precise timing with LSTM recurrent networks. *Journal of Machine Learning Research*, 3, 115–143.
- Goldstein, T., O’Donoghue, B., & Setzer, S. (2014). Fast alternating direction optimization methods. *SIAM Journal on Imaging Sciences*, 7, 1588–1623.
- Golle, P. (2008). Machine learning attacks against the Asirra CAPTCHA. In *ACM CCS*.
- Greff, K., Srivastava, R. K., & Schmidhuber, J. (2016). Highway and residual networks learn unrolled iterative estimation. arXiv preprint [arXiv:1612.07771](https://arxiv.org/abs/1612.07771).
- Gregor, K., & LeCun, Y. (2010). Learning fast approximations of sparse coding. In *ICML*.
- Jarrett, K., Kavukcuoglu, K., Koray, M., & LeCun, Y. (2009). What is the best multi-stage architecture for object recognition? In *ICCV*.
- Kavukcuoglu, K., Ranzato, M. A., & LeCun, Y. (2010). Fast inference in sparse coding algorithms with applications to object recognition. CoRR [arXiv:1010.3467](https://arxiv.org/abs/1010.3467).
- Krizhevsky, A., & Hinton, G. (2009). *Learning multiple layers of features from tiny images* (Vol. 1, No. 4, p. 7). Technical report, University of Toronto.
- Lange, M., Zühlke, D., Holz, O., Villmann, T. (2014). Applications of LP-norms and their smooth approximations for gradient based learning vector quantization. In *ESANN*.
- Le Roux, J., Hershey, J. R., & Weninger, F. (2015). Deep NMF for speech separation. In *ICASSP*.
- LeCun, Y., Bottou, L., Bengio, Y., & Haffner, P. (2009). Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*.
- Liao, Q., & Poggio, T. (2016). Bridging the gaps between residual learning, recurrent neural networks and visual cortex. arXiv preprint [arXiv:1604.03640](https://arxiv.org/abs/1604.03640).
- Liu, S., Xian, Y., Li, H., & Yu, Z. (2017). Text detection in natural scene images using morphological component analysis and laplacian dictionary. *IEEE/CAA Journal of Automatica Sinica*, PP(99), 1–9.
- Moreau, T., & Bruna, J. (2016). Understanding trainable sparse coding with matrix factorization. arXiv preprint [arXiv:1609.00285](https://arxiv.org/abs/1609.00285).
- Nesterov, Y. (2013). *Introductory lectures on convex optimization: A basic course* (Vol. 87). Berlin: Springer.
- Olshausen, B., & Field, D. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381, 607–609.
- Orhan, E., & Pitkow, X. (2018). Skip connections eliminate singularities. In *International conference on learning representations*.
- Otazo, R., Candès, E., & Sodickson, D. K. (2015). Low-rank and sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components. *Magnetic Resonance in Medicine*, 73(3), 1125–36.
- Parekh, A., Selesnick, I., Rapoport, D., & Ayappa, I. (2014). Sleep spindle detection using time-frequency sparsity. In *IEEE SPMB*.
- Peyré, G., Fadili, J., & Starck, J. L. (2007). Learning adapted dictionaries for geometry and texture separation. In *SPIE Wavelets*.
- Peyré, G., Fadili, J., & Starck, J. L. (2010). Learning the morphological diversity. *SIAM Journal on Imaging Sciences*, 3(3), 646–669.
- Schmidt, M., Fung, G., & Rosales, R. (2007). Fast optimization methods for l1 regularization: A comparative study and two new approaches. In J. N. Kok, J. Koronacki, R. L. D. Mantaras, S. Matwin, D. Mladenič, A. Skowron (Eds.), *ECML*.
- Selesnick, I. (2014). L1-norm penalized least squares with salsa. *Connexions* (p. 66). Retrieved March 1, 2017 from <http://cnx.org/contents/e980d3cd-f201-4ef6-8992-d712bf0a88a3@5>.
- Shoham, N., & Elad, M. (2008). Algorithms for signal separation exploiting sparse representations, with application to texture image separation. In *Proceedings of the IEEE 25th convention of electrical and electronics engineers in Israel*.

- Sprechmann, P., Litman, R., Yakar, T., Bronstein, A., & Sapiro, G. (2013). Efficient supervised sparse analysis and synthesis operators. In *NIPS*.
- Starck, J. L., Elad, M., & Donoho, D. (2004). Redundant multiscale transforms and their application for morphological component separation. *Advances in Imaging and Electron Physics*, 132, 287–348.
- Starck, J. L., Elad, M., & Donoho, D. (2005a). Image decomposition via the combination of sparse representations and a variational approach. *IEEE Transactions on Image Processing*, 14(10), 1570–1582.
- Starck, J. L., Moudden, Y., Bobina, J., Elad, M., Donoho, D. (2005b). Morphological component analysis. In *Proceedings of SPIE Wavelets*.
- Tian, S., Pan, Y., Huang, C., Lu, S., Yu, K., & Lim Tan, C. (2015). Text flow: A unified text detection system in natural scene images. In *Proceedings of the IEEE international conference on computer vision* (pp. 4651–4659).
- Uysal, F., Selesnick, I., & Isom, B. (2016). Mitigation of wind turbine clutter for weather radar by signal separation. *IEEE Transactions on Geoscience and Remote Sensing*, 54(5), 2925–2934.
- Wang, Y., Yin, W., & Zeng, J. (2019). Global convergence of ADMM in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, 78(1), 29–63. <https://doi.org/10.1007/s10915-018-0757-z>.
- Wang, Z., Ling, Q., & Huang, T. (2016). Learning deep L0 encoders. In *AAAI*.
- Wisdom, S., Powers, T., Pitton, J., & Atlas, L. (2017). Deep recurrent NMF for speech separation by unfolding iterative thresholding. In *IEEE workshop on applications of signal processing to audio and acoustics (WASPAA)* (pp. 254–258).
- Xiao, H., Rasul, K., & Vollgraf, R. (2017). Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms. CoRR [arXiv:1708.07747](https://arxiv.org/abs/1708.07747).
- Yan, C., Xie, H., Liu, S., Yin, J., Zhang, Y., & Dai, Q. (2018). Effective Uyghur language text detection in complex background images for traffic prompt identification. *IEEE Transactions on Intelligent Transportation Systems*, 19(1), 220–229.
- Yang, Y., Sun, J., Li, H., & Xu, Z. (2016). Deep ADMM-net for compressive sensing MRI. In *NIPS*.
- Zhou, J., Di, K., Du, J., Peng, X., Yang, H., Pan, S.J., Tsang, I. W., Liu, Y., Qin, Z., & Goh, R. (2018). SC2Net: Sparse LSTMs for sparse coding. In *AAAI*.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.