

DICTIONARY LEARNING FROM SPARSELY CORRUPTED OR COMPRESSED SIGNALS

Christoph Studer and Richard G. Baraniuk

Rice University

ABSTRACT

In this paper, we investigate dictionary learning (DL) from sparsely corrupted or compressed signals. We consider three cases: I) the training signals are corrupted, and the locations of the corruptions are known, II) the locations of the sparse corruptions are unknown, and III) DL from compressed measurements, as it occurs in blind compressive sensing. We develop two efficient DL algorithms that are capable of learning dictionaries from sparsely corrupted or compressed measurements. Empirical phase transitions and an in-painting example demonstrate the capabilities of our algorithms.

Index Terms— Dictionary learning, sparse approximation, compressive sensing, signal restoration, in-painting.

1. INTRODUCTION

1.1. Problem statements

We consider learning a dictionary $\mathbf{A} \in \mathbb{C}^{M \times N_a}$ (whose columns have unit Euclidean norm) from a sequence of corrupted M -dimensional measurement vectors

$$\mathbf{z}_i = \mathbf{A}\mathbf{x}_i + \mathbf{B}\mathbf{e}_i, \quad i = 1, \dots, T. \quad (1)$$

Here, $\mathbf{x}_i \in \mathbb{C}^{N_a}$ are sparse vectors, i.e., only a few entries ($n_x \ll N_a$) are non-zero. The dictionary $\mathbf{B} \in \mathbb{C}^{M \times N_b}$ together with the sparse interference vectors $\mathbf{e}_i \in \mathbb{C}^{N_b}$ are used to represent sparse corruptions on the measurements \mathbf{z}_i .

The restoration of \mathbf{x}_i from the sparsely corrupted measurements \mathbf{z}_i in (1), for different cases of knowledge about the support sets of \mathbf{x}_i and \mathbf{e}_i , was investigated previously in [1]. This framework has been shown to feature prominently in numerous applications, such as the restoration of signals from impulse noise, narrowband interference, and saturation, as well as in-painting, super-resolution, or signal separation (see [1] and the references therein).

In addition to learning \mathbf{A} from (1), we are also interested in identifying the dictionary \mathbf{A} from a series of linear and

non-adaptive measurements

$$\tilde{\mathbf{z}}_i = \mathbf{C}_i \mathbf{A} \mathbf{x}_i, \quad i = 1, \dots, T \quad (2)$$

with $\tilde{\mathbf{z}}_i \in \mathbb{C}^{m_i \times 1}$ and $\mathbf{C}_i \in \mathbb{C}^{m_i \times M}$ representing the measurement (or sensing) matrices. If $m_i < M$, the input-output relation (2) corresponds to learning the dictionary \mathbf{A} from a sequence of compressed measurements, as it occurs, for example, in blind compressive sensing (CS) [2, 3].

1.2. Application examples

1) Image in-painting: In-painting [4] amounts to restoring the signal vector $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$ from a subset of its entries; this is captured by (1) by setting $\mathbf{B} = \mathbf{I}_M$ and associating the locations of the nonzero entries of \mathbf{e}_i to the missing pixels in the corrupted image \mathbf{z}_i . In-painting can then be achieved by recovering \mathbf{x}_i from $\mathbf{z}_i = \mathbf{A}\mathbf{x}_i + \mathbf{e}_i$ followed by computing the (full-resolution) signal $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$. The dictionary \mathbf{A} used for in-painting is typically obtained through DL algorithms from a set of uncorrupted (training) images [5].

In [3], the authors proposed to learn the dictionary \mathbf{A} *directly* from the corrupted image. To this end, a novel DL algorithm was developed in [3], which relies on a one-block sparse signal model for the vectors $\mathbf{x}_i, \forall i$ and assumes the locations of the sparse corruptions are known. The framework (1) allows us to generalize [3] to arbitrary (i.e., unstructured) sparse signals. In addition, we will develop a novel algorithm that is agnostic to the support sets of $\mathbf{e}_i, \forall i$.

2) Blind CS: CS aims to acquire a signal \mathbf{y} using fewer measurements than the Nyquist rate suggests [6]. Specifically, CS considers the acquisition of \mathbf{y} through linear and non-adaptive measurements $\tilde{\mathbf{z}} = \mathbf{C}\mathbf{y}$, where \mathbf{C} is a $m \times M$ -dimensional sensing matrix with $m < M$.

In [2], the standard CS framework was extended to *blind CS*, i.e., the *simultaneous* identification of the basis \mathbf{A} and the sparse vectors \mathbf{x}_i from a series of compressed measurements $\tilde{\mathbf{z}}_i, \forall i$ as captured by (2). To establish the uniqueness of blind CS, [2] considered certain constraints on \mathbf{A} . More recently, the blind-CS framework was extended to arbitrary \mathbf{A} and to varying sensing matrices \mathbf{C}_i , but where the signals follow a one-block-sparse model. To overcome these limitations, we develop an algorithm that is able to learn *arbitrary* dictionaries from *general* (i.e., unstructured) sparse signals.

1.3. Contributions

In this paper, we address the question of whether dictionaries can be learned from sparsely corrupted or compressed measurements. To this end, we investigate DL algorithms for (1)

The authors would like to thank E. Dyer, J. Laska, A. Maleki, G. Pope, and A. Sankaranarayanan for inspiring discussions. This work was supported by the Swiss National Science Foundation (SNSF) under Grant PA00P2-134155 and by the Grants NSF CCF-0431150, CCF-0728867, CCF-0926127, DARPA/ONR N66001-08-1-2065, N66001-11-1-4090, N66001-11-C-4092, ONR N00014-08-1-1112, N00014-10-1-0989, AFOSR FA9550-09-1-0432, ARO MURIs W911NF-07-1-0185 and W911NF-09-1-0383, and by the Texas Instruments Leadership University Program. The authors are with Dept. ECE, Rice University, Houston, TX 77005, USA. E-mails: studer@rice.edu; richb@rice.edu. Website: <http://dsp.rice.edu/>

and (2), and study the following cases: I) The locations of the sparse corruptions are unknown, II) the locations of the sparse corruptions are known, and III) the case of blind CS. We develop two efficient DL algorithms: The first is designed specifically for Case I (Sec. 3); the second is able to handle Cases II and III (Sec. 4). Finally, we show empirical phase transitions and an in-painting example to demonstrate the capabilities of the new DL algorithms (Sec. 5).

2. DICTIONARY LEARNING

2.1. The DL problem

For the case $\mathbf{e}_i = \mathbf{0}_{M \times 1}$ and $\mathbf{C}_i = \mathbf{I}_M, \forall i$ we can rewrite (1) and (2) in the more compact form $\mathbf{Z} = \mathbf{A}\mathbf{X}$, with $\mathbf{Z} = [\mathbf{z}_1 \cdots \mathbf{z}_T]$ and $\mathbf{X} = [\mathbf{x}_1 \cdots \mathbf{x}_T]$. Recovery of the dictionary $\mathbf{A} = [\mathbf{a}_1 \cdots \mathbf{a}_{N_a}]$ (along with \mathbf{X}) can be achieved by means of the following DL problem [5]:

$$(\text{DL}) \begin{cases} \text{minimize} & \|\mathbf{Z} - \mathbf{A}\mathbf{X}\|_F^2 \\ \text{subject to} & \|\mathbf{x}_i\|_0 \leq n_x, i = 1, \dots, T \\ & \|\mathbf{a}_\ell\|_2 = 1, \ell = 1, \dots, N_a. \end{cases}$$

Here, $\|\mathbf{x}_i\|_0$ designates the number of non-zeros in \mathbf{x}_i and $\|\cdot\|_F$ the Frobenius norm. Solving the DL problem requires combinatorial complexity, in general, and, hence, a variety of computationally efficient approximations have been proposed. Prominent candidate algorithms are maximum-likelihood-based methods, maximum a posteriori (or Bayesian) methods, the method of optimal directions, and the K-SVD algorithm (see [5] and the references therein).

2.2. The K-SVD algorithm

Since the DL algorithms developed in the remainder of the paper build upon on ideas of the K-SVD algorithm [5], we briefly review its main steps. The K-SVD algorithm is initialized by a (random or well-defined) dictionary \mathbf{A} and then repeats the following two steps to reduce the objective of the DL problem. The iterative procedure is terminated if either the objective function is below a certain threshold or a maximum number of iterations has been reached.

1) *Sparse representations*: In this step, \mathbf{A} is held constant and sparse representations $\hat{\mathbf{x}}_i$ for each observation vector \mathbf{z}_i are computed. One is typically interested in an approximate solution $\hat{\mathbf{x}}_i$ to the *primal sparse approximation problem* [7]

$$(\text{PSAP}) \text{ minimize } \|\mathbf{z}_i - \mathbf{A}\mathbf{x}_i\|_2^2 \text{ s.t. } \|\mathbf{x}_i\|_0 \leq n_x,$$

which can, for example, be achieved by means of orthogonal matching pursuit (OMP) [7] (or other greedy pursuits).

2) *Updating atoms and coefficients*: In this step, the atoms $\mathbf{a}_\ell, \ell = 1, \dots, N_a$ along with the non-zero entries in \mathbf{X} , are updated sequentially. To this end, the ℓ th atom \mathbf{a}_ℓ is isolated as

$$\|\mathbf{Z} - \mathbf{A}\mathbf{X}\|_F^2 = \|\mathbf{S}_\ell - \mathbf{a}_\ell \bar{\mathbf{x}}_\ell^H\|_F^2 \quad (3)$$

where the row-vector $\bar{\mathbf{x}}_\ell^H$ corresponds to the ℓ th row of the matrix \mathbf{X} and $\mathbf{S}_\ell = \mathbf{Z} - \sum_{j \neq \ell} \mathbf{a}_j \bar{\mathbf{x}}_j^H$. To minimize (3), the

algorithm now *simultaneously* updates the atom \mathbf{a}_ℓ and the non-zero entries in $\bar{\mathbf{x}}_\ell^H$. This can be achieved by defining Ω_ℓ as the index set of the non-zero entries in $\bar{\mathbf{x}}_\ell^H$ and by minimizing $\|\mathbf{S}_\ell\|_{\Omega_\ell} - \mathbf{a}_\ell [\bar{\mathbf{x}}_\ell^H]_{\Omega_\ell}\|_F^2$, where $[\mathbf{S}_\ell]_{\Omega_\ell} \in \mathbb{C}^{M \times |\Omega_\ell|}$ and $[\bar{\mathbf{x}}_\ell^H]_{\Omega_\ell} \in \mathbb{C}^{1 \times |\Omega_\ell|}$ consist of those columns of \mathbf{S}_ℓ and entries of $\bar{\mathbf{x}}_\ell^H$ belonging to Ω_ℓ . Finding \mathbf{a}_ℓ and $[\bar{\mathbf{x}}_\ell^H]_{\Omega_\ell}$ that minimize the Frobenius norm amounts to computing the best rank-1 approximation to $[\mathbf{S}_\ell]_{\Omega_\ell}$; this is commonly achieved by using the singular value decomposition (SVD) of $[\mathbf{S}_\ell]_{\Omega_\ell}$ [5].

3. DL FOR UNKNOWN INTERFERENCE SUPPORT

In this section, we develop a DL algorithm for Case I, i.e., we are interested in retrieving \mathbf{A} and the sparse vectors $\mathbf{x}_i, \forall i$ from (1), where no knowledge about the support of $\mathbf{e}_i, \forall i$, apart from its cardinality n_e , is available; the dictionary \mathbf{B} is assumed to be known.

3.1. The partial DL problem

Analogously to the DL problem in Sec. 2.1, we define the partial DL problem (P-DL) for Case I as follows:

$$(\text{P-DL}) \begin{cases} \text{minimize} & \|\mathbf{Z} - \mathbf{A}\mathbf{X} - \mathbf{B}\mathbf{E}\|_F^2 \\ \text{subject to} & \|\mathbf{x}_i\|_0 \leq n_x, i = 1, \dots, T \\ & \|\mathbf{e}_i\|_0 \leq n_e, i = 1, \dots, T \\ & \|\mathbf{a}_\ell\|_2 = 1, \ell = 1, \dots, N_a \end{cases}$$

with $\mathbf{E} = [\mathbf{e}_1 \cdots \mathbf{e}_T]$. Solving P-DL requires combinatorial complexity, which necessitates efficient approximations. The algorithm described next extends the K-SVD algorithm [5] and delivers an approximation to the P-DL problem.

3.2. The partial K-SVD (PK-SVD) algorithm

The dictionary \mathbf{A} is initialized as for the K-SVD algorithm. The following two steps are repeated either until the objective function of the P-DL problem is below a certain threshold or a maximum number of iterations has been reached.

1) *Sparse representations*: Since we assume knowledge of both n_x and n_e in the P-DL problem, we intend to find $\forall i$ the solutions $\hat{\mathbf{x}}_i$ and $\hat{\mathbf{e}}_i$ of the separation PSAP:

$$(\text{S-PSAP}) \begin{cases} \text{minimize} & \|\mathbf{z}_i - \mathbf{A}\mathbf{x}_i - \mathbf{B}\mathbf{e}_i\|_2^2 \\ \text{subject to} & \|\mathbf{x}_i\|_0 \leq n_x \text{ and } \|\mathbf{e}_i\|_0 \leq n_e. \end{cases}$$

To efficiently obtain an approximate solution for S-PSAP, we propose to use a modified version of OMP. Specifically, we limit the maximum number of atoms allowed to be chosen from \mathbf{A} and \mathbf{B} by n_x and n_e , respectively, which ensures that the so-obtained vectors satisfy $\|\mathbf{x}_i\|_0 \leq n_x$ and $\|\mathbf{e}_i\|_0 \leq n_e$.

2) *Updating atoms and coefficients*: In this step, we only update the atoms in \mathbf{A} using the same procedure as for the K-SVD; both the matrices \mathbf{B} and \mathbf{E} are held constant.

4. DL FOR KNOWN INTERFERENCE SUPPORT OR COMPRESSED MEASUREMENTS

In this section, we develop a DL algorithm for both Cases II and III. For Case II, the support sets \mathcal{E}_i of each interference

vector \mathbf{e}_i is assumed to be known; this allows us to project \mathbf{z}_i onto the orthogonal complement of the range space spanned by $\mathbf{B}_{\mathcal{E}_i}$, which leads to the following input-output relation [1]:

$$\tilde{\mathbf{z}}_i = \mathbf{R}_{\mathcal{E}_i} \mathbf{z}_i = \mathbf{R}_{\mathcal{E}_i} \mathbf{A} \mathbf{x}_i, \quad i = 1, \dots, T. \quad (4)$$

Here, $\mathbf{R}_{\mathcal{E}_i} = \mathbf{I}_M - \mathbf{B}_{\mathcal{E}_i} \mathbf{B}_{\mathcal{E}_i}^\dagger$, and we used $\mathbf{R}_{\mathcal{E}_i} \mathbf{B}_{\mathcal{E}_i} = \mathbf{0}_{M \times 1}$. The DL algorithm presented next will be formulated for the more general Case III, i.e., for the input-output relation (2), as opposed to the special case in (4), for which $\mathbf{C}_i = \mathbf{R}_{\mathcal{E}_i}$.

4.1. The subspace DL problem

The DL problem for (2), referred to as subspace DL (S-DL) in the remainder of the paper, aims to identify \mathbf{A} (together with the coefficients $\mathbf{x}_i, \forall i$) as

$$\text{(S-DL)} \quad \begin{cases} \text{minimize} & \sum_{i=1}^T \|\tilde{\mathbf{z}}_i - \mathbf{C}_i \mathbf{A} \mathbf{x}_i\|_2^2 \\ \text{subject to} & \|\mathbf{x}_i\|_0 \leq n_x, i = 1, \dots, T \\ & \|\mathbf{a}_\ell\|_2 = 1, \ell = 1, \dots, N_a. \end{cases}$$

It is important to realize that for $\mathbf{C}_i \neq \mathbf{C}_j$, for some $i \neq j$, or for non-orthogonal \mathbf{C}_i , one cannot rewrite the objective of S-DL as in (3); this property inhibits the use of the K-SVD algorithm. A suitable DL algorithm is described next.

4.2. The subspace-DL algorithm

The algorithm developed next is inspired by the K-SVD algorithm and referred to as subspace DL-algorithm (SDLA) in the following. The dictionary \mathbf{A} is initialized as for the K-SVD algorithm. Then, we iteratively reduce the objective of the S-DL problem, by performing the following two steps until either the objective is below a certain threshold or a maximum number of iterations has been reached.

1) *Sparse representations*: The first step computes sparse representations $\hat{\mathbf{x}}_i$ for each measurement \mathbf{z}_i as the solution of the following subspace PSAP:

$$\text{(S-PSAP)} \quad \text{minimize } \|\tilde{\mathbf{z}}_i - \mathbf{C}_i \mathbf{A} \mathbf{x}_i\|_2^2 \quad \text{s.t. } \|\mathbf{x}_i\|_0 \leq n_x.$$

The solution of this problem can be approximated using OMP.

2) *Updating atoms and coefficients*: We start by rewriting the objective of the S-DL problem as

$$\sum_{i=1}^T \|\hat{\mathbf{z}}_i^{(\ell)} - \mathbf{C}_i \mathbf{a}_\ell [\mathbf{x}_i]_\ell\|_2^2, \quad (5)$$

where $\hat{\mathbf{z}}_i^{(\ell)} = \tilde{\mathbf{z}}_i - \sum_{j \neq \ell} \mathbf{C}_i \mathbf{a}_j [\mathbf{x}_i]_j$. It is important to realize that (5) cannot be optimized using rank-1 approximations. A common approach is to employ alternating optimization, i.e., to alternate between atom and coefficient updates. Following this approach, updating the ℓ -th atom \mathbf{a}_ℓ amounts to solving

$$\hat{\mathbf{a}}_\ell = \arg \min_{\mathbf{a}_\ell, \|\mathbf{a}_\ell\|_2=1} \sum_{i=1}^T \|\hat{\mathbf{z}}_i^{(\ell)} - \mathbf{C}_i \mathbf{a}_\ell [\mathbf{x}_i]_\ell\|_2^2 \quad (6)$$

for which iterative optimization techniques exist. To avoid the computation of (6) and alternate optimization altogether, let $\mathbf{p}_\ell = \mathbf{a}_\ell$ and consider the relaxation of (6) to

$$\hat{\mathbf{p}}_\ell = \arg \min_{\mathbf{p}_\ell \in \mathbb{C}^{M \times N_a}} \sum_{i=1}^T \|\hat{\mathbf{z}}_i^{(\ell)} - \mathbf{C}_i \mathbf{p}_\ell [\mathbf{x}_i]_\ell\|_2^2. \quad (7)$$

Omitting the unit-norm constraint in (7) enables us to obtain $\hat{\mathbf{p}}_\ell$ in closed form. To this end, we set the partial derivative $\partial / (\partial \mathbf{p}_\ell^H)$ of the objective in (7) to zero, leading to

$$\sum_{i=1}^T \left(\|\mathbf{x}_i\|_\ell^2 \mathbf{C}_i^H \mathbf{C}_i \mathbf{p}_\ell - [\mathbf{x}_i]_\ell^* \mathbf{C}_i^H \hat{\mathbf{z}}_i^{(\ell)} \right) = \mathbf{0}_{M \times 1},$$

which leads to the following closed-form solution:

$$\hat{\mathbf{p}}_\ell = \left(\sum_{i=1}^T \|\mathbf{x}_i\|_\ell^2 \mathbf{C}_i^H \mathbf{C}_i \right)^\dagger \sum_{i=1}^T [\mathbf{x}_i]_\ell^* \mathbf{C}_i^H \hat{\mathbf{z}}_i^{(\ell)}.$$

The pseudo-inverse \mathbf{M}^\dagger is defined through the SVD of $\mathbf{M} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^H$ as $\mathbf{M}^\dagger = \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}^H$ with the diagonal matrix $\mathbf{\Sigma}^{-1}$ defined as $[\mathbf{\Sigma}^{-1}]_{k,k} = [\mathbf{\Sigma}]_{k,k}^{-1}$ if $[\mathbf{\Sigma}]_{k,k} > 0$ and $[\mathbf{\Sigma}^{-1}]_{k,k} = 0$ otherwise; this ensures that $\hat{\mathbf{p}}_\ell$ can be computed in any case.

Even though the vectors $\hat{\mathbf{p}}_\ell$ do not necessarily have unit ℓ_2 -norm, we can absorb this constraint into $[\mathbf{x}_i]_\ell, \forall i$, as

$$\hat{\mathbf{a}}_\ell = \hat{\mathbf{p}}_\ell / \|\hat{\mathbf{p}}_\ell\|_2 \quad \text{and} \quad [\hat{\mathbf{x}}_i]_\ell = [\mathbf{x}_i]_\ell \|\hat{\mathbf{p}}_\ell\|_2, \quad i = 1, \dots, T,$$

which does not affect the objective in (7) and ensures that $\|\hat{\mathbf{a}}_\ell\|_2 = 1, \forall \ell$. We emphasize that this atom and coefficient update procedure results in significantly smaller complexity compared to an alternate-optimization-based approach.

5. NUMERICAL RESULTS

5.1. Empirical phase transitions

We next illustrate the impact of support-set knowledge on the DL algorithms detailed above and provide a comparison to the K-SVD algorithm [5], for which no corruptions are assumed.

1) *Simulation procedure*: For each pair $(\mathbf{x}_i, \mathbf{e}_i)$ of (n_x, n_e) -sparse vectors, 10 Monte-Carlo trials are performed. The support sets are generated uniformly at random and the non-zero entries are i.i.d. zero-mean Gaussian and $T = 2000$. The dictionary \mathbf{A} to be retrieved corresponds to a 32×48 over-complete DCT dictionary; the interference is assumed to be sparse in the identity, i.e., $\mathbf{B} = \mathbf{I}_{32}$. Figure 1 show the average (over all N_a atoms) success rate; success is declared if $|\mathbf{a}_\ell^H \tilde{\mathbf{a}}_\ell| \geq 0.95$ with $\tilde{\mathbf{a}}_\ell$ denoting the recovered atom.

2) *Discussion*: Figure 1 shows that the K-SVD is able to retrieve \mathbf{A} for up to $n_x = 9$. For the SDLA, we see that corrupted measurements degrade the probability of success roughly linear in n_e . For a given number n_e , the difference of SDLA Fig. 1(b) to the PK-SVD algorithm Fig. 1(c) is that knowing the interference support allows for approximately twice the number of corruptions in the measurements. We emphasize that this factor-of-two penalty can be observed in the analytical recovery guarantees of sparsely corrupted signals (see [1] for the details).

5.2. Image in-painting

We now show an in-painting example, where we learn the dictionary \mathbf{A} directly from the corrupted image. We furthermore provide a comparison to the blind-DL algorithm (BDLA) in [3] specifically designed for one-block sparse signals. Our goal here is not to benchmark our performance vs. well-established in-painting methods (e.g., [4]), but rather

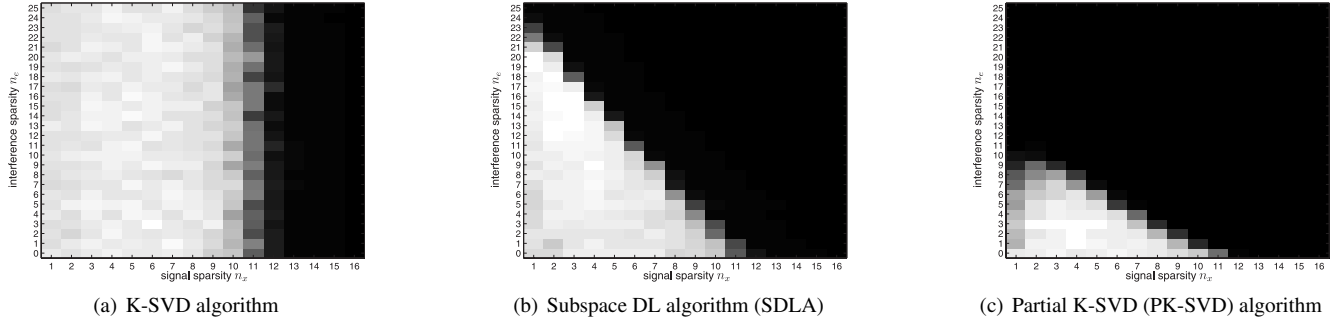


Fig. 1. Empirical phase transitions of K-SVD, the SDLA, and PK-SVD (white corresponds to 100% success of recovery).



Fig. 2. In-painting results using SDLA and BDLA [3] for an image with 50% missing entries.

to demonstrate that in-painting can be performed with dictionaries learned from the corrupted image itself using SDLA.

1) *Corruption and in-painting procedure:* We corrupt 50% of the entries of a 512×512 image by setting the corrupted entries to zero (the pixel intensities are within $[-1, 1]$). DL is performed as follows: We learn a 64×128 dictionary \mathbf{A} by the SDLA and the BDLA [3] from 20 000 randomly selected 8×8 patches. 20 (outer) iterations are performed and we recover signals having 8 non-zero entries. Finally, in-painting is performed by a simple weighted averaging of $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$ over all training patches. We use the reconstruction signal-to-noise-ratio (RSNR) as performance measure.

2) *Discussion:* Figure 2 shows the results of the in-painting procedure. The image recovered by the SDLA has an RSNR of 18.8 dB and only required 27 min. on a 2.3 GHz mobile quad-core CPU. For the BDLA [3], we obtain 14.0 dB RSNR while requiring 658 min. We conclude the following: The SDLA i) achieves better RSNR than BDLA [3] as it is not restricted to a one-block-sparse signal model, ii) exhibits significantly lower computational complexity, and iii) has a simple structure, whereas BDLA [3] is a combination of several (computationally and structurally complex) techniques.

6. CONCLUSION

We have developed two DL algorithms for signals that are corrupted by sparse interference or obtained through compressed measurements. Our DL algorithms are applicable to a variety of signal-restoration applications [1] and to blind CS [2, 3]. The presented numerical experiments have shown

that DL from corrupted or compressed measurements is feasible with computationally efficient algorithms.

7. REFERENCES

- [1] C. Studer, P. Kuppinger, G. Pope, and H. Bölcskei, “Recovery of sparsely corrupted signals,” *to appear in IEEE Trans. Inf. Theory*, 2012.
- [2] S. Gleichman and Y. C. Eldar, “Blind compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 57, no. 10, pp. 6958–6975, Oct. 2011.
- [3] J. Silva, M. Chen, Y. C. Eldar, G. Sapiro, and L. Carin, “Blind compressed sensing over a structured union of subspaces,” *arXiv:1103.2469*, Mar. 2011.
- [4] M. Bertalmio, G. Sapiro, V. Caselles, and C. Ballester, “Image inpainting,” *Proc. SIGGRAPH*, pp. 417–424, 2000.
- [5] M. Aharon, M. Elad, and A. M. Bruckstein, “K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation,” *IEEE Trans. Sig. Proc.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [6] D. L. Donoho, “Compressed sensing,” *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [7] J. A. Tropp, “Greed is good: Algorithmic results for sparse approximation,” *IEEE Trans. Inf. Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.