King's Research Portal

DOI:

[10.1109/TSP.2020.2971948](https://doi.org/10.1109/TSP.2020.2971948)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Yu, Q., Dai, W., Cvetkovic, Z., & Zhu, J. (2020). Dictionary Learning with BLOTLESS Update. *IEEE Transactions on Signal Processing*, 68, 1635-1645. [8985423]. <https://doi.org/10.1109/TSP.2020.2971948>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

Dictionary Learning with BLOTLESS Update

Qi Yu^{*†}, Wei Dai[†], Zoran Cvetkovic[‡], Jubo Zhu^{*}

^{*}College of Liberal Arts and Sciences, National University of Defense Technology, Changsha, China

[†]Department of Electrical and Electronic Engineering, Imperial College London, UK

[‡]Department of Informatics, Kings College London, UK

Abstract—Algorithms for learning a dictionary to sparsely represent a given dataset typically alternate between sparse coding and dictionary update stages. Methods for dictionary update aim to minimise expansion error by updating dictionary vectors and expansion coefficients given patterns of non-zero coefficients obtained in the sparse coding stage. We propose a block total least squares (BLOTLESS) algorithm for dictionary update. BLOTLESS updates a block of dictionary elements and the corresponding sparse coefficients simultaneously. In the error free case, three necessary conditions for exact recovery are identified. Lower bounds on the number of training data are established so that the necessary conditions hold with high probability. Numerical simulations show that the bounds approximate well the number of training data needed for exact dictionary recovery. Numerical experiments further demonstrate several benefits of dictionary learning with BLOTLESS update compared with state-of-the-art algorithms especially when the amount of training data is small.

I. INTRODUCTION

Sparse signal representation has found a wide range of applications, including image denoising [1], [2], image inpainting [1], image deconvolution [3], image super-resolution [4], [5], etc. The key idea behind the concept of sparse representation is that natural signals tend to have sparse representations under certain bases/dictionaries. Hence, finding a dictionary under which a given data set can be represented in a sparse manner, has become a very active area of research. Although numerous analytical dictionaries exist, including Fourier basis [6], discrete cosine transform (DCT) dictionaries, wavelets [7], curvelets [8], etc., the need to adapt to properties of specific data sets has long been driving research efforts towards the development of efficient algorithms for dictionary learning [9], [10]. More formally, dictionary learning is the problem of finding a dictionary $\mathbf{D} \in \mathbb{R}^{m \times l}$ of l vectors in $\mathbb{R}^m$ such that the n training samples in $\mathbf{Y} \in \mathbb{R}^{m \times n}$ can be written as $\mathbf{Y} = \mathbf{DX}$, where the coefficient matrix $\mathbf{X} \in \mathbb{R}^{l \times n}$ is sparse. Of particular interest is overcomplete dictionary learning where the number of dictionary items is larger than the data dimension, i.e., $l > m$, and the number of the training

samples is typically much larger than the size of the dictionary, $n \gg l$. Dictionary learning is a nonconvex bilinear inverse problem, very challenging to solve in general.

The bilinear dictionary learning problem is typically approached by alternating between two stages: sparse coding and dictionary update [9]–[14]. In the sparse coding stage, the goal is to find sparse representations $\mathbf{X}$ of training samples $\mathbf{Y}$ for a given dictionary $\mathbf{D}$. For that purpose, scores of algorithms have been developed. They can be divided into two main categories. The first category consists of greedy algorithms, including orthogonal matching pursuit (OMP) [15], regularized orthogonal matching pursuit (ROMP) [16], subspace pursuit (SP) [17], etc. In the second category, sparse coding is formulated as a convex optimization problem where ℓ_1 -norm is used to promote sparsity [18], and then optimization techniques, e.g. the fast iterative shrinkage-thresholding algorithm (FISTA) [19], can be applied. Reviews of sparse recovery algorithms can be found in [20].

The goal of the dictionary update is to refine the dictionary so that the training samples $\mathbf{Y}$ have more accurate sparse representations given indices of non-zero coefficients obtained in the sparse coding stage. In the probabilistic framework, one may apply either maximum likelihood (ML) estimator [9] or maximum a posteriori (MAP) estimator [12], and then solve them by using gradient decent procedures. In the context of ML formulation [9], Engan *et al.* [11] proposed the method of optimal directions (MOD) where the sparse coefficients $\mathbf{X}$ are fixed and the dictionary update problem is cast as a least squares problem which can be solved efficiently; modifications of MOD were subsequently proposed in [21]–[23].

Recently an alternative approach for dictionary update has become popular, where both the dictionary and the sparse coefficients are updated *simultaneously* with a given sparsity pattern. The representative algorithms include the famous K-SVD algorithm [10], [24] and SimCO [13]. The crux of K-SVD [10] algorithm is to update dictionary items and their corresponding sparse coefficients simultaneously, sequentially one by one. K-SVD was subsequently extended to allow simultaneous update of multiple dictionary elements and corresponding coefficients [24]. SimCO [13], of which K-SVD is a special case, goes further and updates the whole dictionary and sparse coefficients simultaneously. The main idea of SimCO is that given a sparsity pattern, the sparse coefficients can be viewed as a function of the dictionary. As a result, the dictionary update becomes a nonconvex optimisation problem with respect to the dictionary. The optimisation is then preformed using the gradient descent method combined with a heuristic

This work is supported by the National Key Research and Development Program of China 2017YFB0502703, the National Natural Science Foundation of China under Grant 61671015, China Scholarship Council and Royal Society International Exchanges 2017 Cost Share (with China). The conference version of this paper was presented at the International Conference on Acoustics, Speech, and Signal Processing (ICASSP), Brighton, UK, May 2019.

The authors also thank suggestions at ICASSP 2019 on structured total least squares given by Prof. Yoram Bresler from University of Illinois Urbana-Champaign and Dr. Konstantin Usevich from French National Center for Scientific Research.

sub-routine designed to deal with singular points which can prevent from the convergence to the global minimum [13].

Due to the non-convexity of dictionary learning problem, it is challenging to understand under which conditions exact dictionary recovery is possible and which method is optimal in achieving that. Following early efforts on theoretical analysis of exact dictionary recovery [25]–[31], more recently, Spielman et. al. [32] studied dictionary learning problem with complete dictionaries where the dictionary can be presented as a square matrix. By solving a certain sequence of linear programs, they showed that one can recover a complete dictionary D from $Y = DX$ when X is a sparse random matrix with $O(\sqrt{m})$ nonzeros per column. In [33]–[36], the authors propose algorithms which combine clustering, spectral initialization, and local refinement to recover overcomplete and incoherent dictionaries.

Again these algorithms succeed when X has $O(\sqrt{m})$ nonzeros per column. The work in [37] provides a polynomial-time algorithm that recovers a large class of over-complete dictionaries when X has $O(m^{1-\delta})$ nonzeros per column for any constant $\delta \in (0, 1)$. However, the proposed algorithm runs in super-polynomial time when the sparsity level goes up to $O(m)$. Similarly, in [38] the authors proposed a super-polynomial time algorithm that guarantees recovery with close to $O(m)$ nonzeros per column. Sun et al. [39], [40], on the other hand, proposed a polynomial-time algorithm that provably recovers complete dictionary D when X has $O(m)$ nonzeros per column and the size of training samples is $O(m^2 \log(m))$.

This paper addresses the dictionary update problem, where both the dictionary and the sparse coefficients are updated, for a given sparsity pattern. Whilst it is a sub-problem of the overall dictionary learning, it is nevertheless an important step towards solving the overall problem, and its bilinear nature makes it nonconvex and hence very challenging to solve. Our main contributions are as follows.

- BLOTLESS simultaneously updates a block of dictionary items and the corresponding sparse coefficients. Inspired by ideas presented in [41], [42], in BLOTLESS the bilinear nonconvex block update problem is transformed into a linear least squares problem, which can be solved efficiently.
- For the error-free case, when the sparsity pattern is known exactly, three necessary conditions for unique recovery are identified, expressed in terms of lower bounds on the number of training data. Numerical simulations show that the theoretical bounds well approximate the empirical number of training data needed for exact dictionary recovery. In particular, we show that the number of training samples needed is $O(m)$ for complete dictionary update.
- BLOTLESS is numerically demonstrated robust to errors in the assumed sparsity pattern. When embedded into the overall dictionary learning process, it leads to faster convergence rate and less training samples needed compared to state-of-the-art algorithms including MOD, K-SVD and SimCO.

Our work is inspired by a recent work [41] where bilinear inverse problems are formulated as linear inverse problems. The main difference is that our theoretical analysis and algorithm designs in Sections III and IV are specifically tailored

to the generic dictionary update problem while the focus of [41] is self-calibration which can be viewed as dictionary learning with only diagonal dictionaries. Parts of the results in this paper were presented in the conference paper [43]. In this journal paper, we refine the bounds in Section III and provide detailed proofs, add two total least squares algorithms in Section IV, and include more simulation results in Section V to support the design of the algorithm.

This paper is organized as follows. Section II briefly reviews dictionary learning and update methods. Section III discusses an ideal case where exact dictionary recovery is possible, for which a least squares method is developed and analysed. In Section IV, the general case of dictionary update is discussed, and the least squares method is extended to total least squares methods, leading to BLOTLESS. Results of extensive simulations are presented in Section V and conclusions are drawn in Section VI.

A. Notation

In this paper, $\|\cdot\|_2$ denotes the ℓ_2 norm and $\|\cdot\|_F$ stands for the Frobenius norm. For a positive integer n , define $[n] = \{1, 2, \dots, n\}$. For a matrix M , $M_{i,:}$ and $M_{:,j}$ denote the i -th row and the j -th column of M respectively. Consider the sparse coefficient matrix X . Let Ω be the support set of X , i.e., the index set that containing indices of all nonzero entries in X . Let Ω_i be the support set of the row vector $X_{i,:}$. Then X_{i,Ω_i} is the row vector obtained by keeping the nonzero entries of $X_{i,:}$ and removing all its zero entries. Symbols I , $\mathbf{1}$, and $\mathbf{0}$ denote the identity matrix, the vector of which all the entries are 1, and the vector with all zero entries, respectively. For a given set $\Omega_i \subset [n]$, Ω_i^c denotes its complement in $[n]$.

II. DICTIONARY LEARNING: THE BACKGROUND

Dictionary learning is the process of finding a dictionary which sparsely represents given training samples. Let $Y \in \mathbb{R}^{m \times n}$ be the training sample matrix, where m is the dimension of training sample vectors and n is the number of training samples. The overall dictionary learning problem is often formulated as:

$$\min_{D, X} \|Y - DX\|_F^2, \text{ s.t. } \|X_{:,j}\|_0 \leq k, \forall j \in [n], \quad (1)$$

where $D \in \mathbb{R}^{m \times l}$ is the dictionary, $X \in \mathbb{R}^{l \times n}$ is the sparse coefficient matrix, the ℓ_0 pseudo-norm $\|\cdot\|_0$ gives the number of non-zero elements, also known as sparsity level, and $k < l$ is the upper bound of the sparsity level.

Dictionary learning algorithms typically iterate between two stages: sparse coding and dictionary update. The goal of sparsity coding is to find a sparse coefficient matrix X for a given dictionary D . One way to achieve this is to solve the problem

$$\min_{X_{:,j}} \|Y_{:,j} - DX_{:,j}\|_2, \text{ s.t. } \|X_{:,j}\|_0 \leq k, \forall j \in [n], \quad (2)$$

using greedy algorithms, for example OMP [15] or SP [17].

In the dictionary update stage, the goal is to refine the dictionary with either fixed sparse coefficients or a fixed sparse pattern, i.e. fixed locations of non-zero coefficients. The

famous MOD method [11] falls into the first category. With fixed sparse coefficients, dictionary update is simply a least squares problem

$$\min_D \|Y - DX\|_F^2.$$

A more popular and advantageous approach is to simultaneously update the dictionary and nonzero sparse coefficients by fixing only the sparsity pattern. With this idea, dictionary update is then formulated as [10], [13], [24]

$$\min_{D, X} \|Y - DX\|_F^2, \text{ s.t. } \mathcal{P}_{\Omega^c}(X) = 0 \quad (3)$$

where $\mathcal{P}_{\Omega^c}(X)$ gives the vector formed by the entries of X indexed by Ω^c . However, problem (3) is bilinear, nonconvex, and challenging to solve.

Among many methods for solving (3), we here briefly review K-SVD [10] and SimCO [13]. K-SVD algorithm successively updates individual dictionary items $D_{:,i}$ and the corresponding sparse coefficients $X_{i,:}$ whilst keeping all other dictionary items and coefficients fixed:

$$\min_{D_{:,i}, X_{i,:}} \|(Y - D_{:, \{i\}^c} X_{\{i\}^c, :})_{:, \Omega_i} - D_{:,i} X_{i, \Omega_i}\|_F^2. \quad (4)$$

The optimal solution can be obtained by taking the largest left and right singular vectors of the matrix $(Y - D_{:, \{i\}^c} X_{\{i\}^c, :})_{:, \Omega_i}$.

The idea of SimCO is to formulate the dictionary update problem in (3) as a nonconvex optimisation problem with respect to the overall dictionary, that is

$$\min_D \underbrace{\min_{X: \mathcal{P}_{\Omega^c}(X)=0} \|Y - DX\|_F^2}_{f(D)}. \quad (5)$$

and then solve it using gradient descent of D . This leads to an update of all dictionary items and sparse coefficients simultaneously. K-SVD can be viewed as a special case of SimCO where the objective function reads

$$\min_{D_{:,i}} \min_{X_{i,:}: \mathcal{P}_{\Omega_i^c}(X_{i,:})=0} \|Y - DX\|_F^2.$$

The focus of this paper is a novel solution to Problem (3).

III. EXACT DICTIONARY RECOVERY

This section focuses on an ideal case that the dictionary can be exactly recovered. We assume that the training samples in Y are generated from $Y = D_0 X_0$ where D_0 is a tall or square matrix ($m \geq l$) and the sparsity pattern of X_0 is given. For compositional convenience, we focus on the case where D_0 is a square matrix, $D_0 \in \mathbb{R}^{m \times m}$, as the same analysis is valid for a tall dictionary where $m > l$.

With given sparsity pattern denoted by Ω , the dictionary update problem can be formulated as a bilinear inverse problem in which the goal is to find D and X such that

$$Y = DX \text{ and } \mathcal{P}_{\Omega^c}(X) = 0. \quad (6)$$

The constraint $Y = DX$ is nonconvex. Generally speaking, it is challenging to solve (6) and there are no guarantees for the global optimality of the solution.

1) *Least Squares Solver*: Suppose that the unknown dictionary matrix D is invertible. The nonconvex problem in (6) can be translated into a convex problem by using a strategy similar to that explored in [41]. Define $H = D^{-1}$. Then $HY = X$. The goal is now to find H and X such that

$$HY = X \text{ and } \mathcal{P}_{\Omega^c}(X) = 0, \quad (7)$$

or equivalently,

$$\begin{bmatrix} Y^T & -I_{n \times n} \end{bmatrix} \begin{bmatrix} H^T \\ X^T \end{bmatrix} = \mathbf{0}_{n \times m} \text{ and } \mathcal{P}_{\Omega^c}(X) = 0, \quad (8)$$

where the subscripts are used to indicate matrix dimensions. In this manner the original bilinear problem (6) is cast as an equivalent linear least squares problem.

However, the formulation in (8) admits trivial solution $H = 0$ and $X = 0$. In fact, (8) admits at least m linearly independent solutions.

Proposition 1. *There are at least m linear independent solutions to the least squares problem in (8).*

Proof. This proposition is proved by construction. Let $H_0 = D_0^{-1}$. Define matrix $Z_i \in \mathbb{R}^{(m+n) \times m}$ by keeping the i -th column of the matrix $[H_0, X_0]^T$ and setting all other columns to zero, that is, $(Z_i)_{:,i} = [(H_0)_{:,i}, (X_0)_{i,:}]^T$ and $(Z_i)_{:,j} = 0$ for all $j \neq i$. From the fact that $(H_0)_{:,i} Y = (X_0)_{i,:}$, it is straightforward to verify that $Z_i, i \in [m]$, is a solution of (8).

The solutions $Z_i, i \in [m]$, are linearly independent. This can be easily verified by observing that the positions of nonzero elements in Z_i and $Z_j, i \neq j$, are different. $\square$

2) *Necessary Conditions for Unique Recovery* : We now consider the uniqueness of the solution in more detail and derive necessary conditions for unique recovery. Two ambiguities can be identified in the dictionary update problem in (8). The first is permutation ambiguity. Let Ω_i and Ω_j be the support sets (the index set containing indices corresponding to nonzero entries) of the i -th and j -th row of X_0 . If $\Omega_i = \Omega_j$, then the tuple $(D_0 P_{i \leftrightarrow j}, P_{i \leftrightarrow j} X_0)$ is a valid solution of (6), where $P_{i \leftrightarrow j}$ denotes the permutation matrix generated by permuting the i -th and j -th row of the identity matrix. On the other hand, there is no permutation ambiguity if $\Omega_i \neq \Omega_j$ for all $i \neq j$. In practice, the given sparsity pattern is typically diverse enough to avoid permutation ambiguity.

The second is the scaling ambiguity which cannot be avoided. Let S be a diagonal matrix with nonzero diagonal elements. It is clear that the tuple $(D_0 S, S^{-1} X_0)$ is also a valid solution of (6). All solutions of the form $(D_0 S, S^{-1} X_0)$ form an equivalent class. The scaling ambiguity can be addressed by introducing additional constraints. One option used in [41] is that the sum of the elements in each row of H is one, i.e., $\sum_j H_{i,j} = 1, \forall i$. With these constraints, one has

$$H[Y, \mathbf{1}_{m \times 1}] = [X, \mathbf{1}_{m \times 1}] \text{ and } \mathcal{P}_{\Omega^c}(X) = 0. \quad (9)$$

Henceforth, we define unique recovery as unique up to the scaling ambiguity.

Definition 1 (Unique Recovery). The dictionary update problem (6) is said to admit a unique solution if all solutions are

of the form $D = D_0 S$ and $X = S^{-1} X_0$ for some diagonal matrix S with nonzero diagonal elements.

In the following, we identified three necessary conditions for unique recovery.

Proposition 2. Assume that D_0 is square and invertible. If the problem (6) has unique solution, then it holds that

- 1) $n \geq n_0 = m + \frac{|\Omega|}{m} - 1$.
- 2) For all $i \in [m]$, the support set of the i -th row of X_0 , denoted by Ω_i , satisfies $|\Omega_i^c| \geq m - 1$.
- 3) For all $i \in [m]$ and all $i' \neq i$, $\exists j \in [n]$ such that $(X_0)_{i,j} = 0$ and $(X_0)_{i',j} \neq 0$.

Proof. Necessary condition 1 is proved by using the fact that the solution of (9) is unique only if the number of equations is larger or equal than the number of unknown variables. The number of unknown variables in (9) is $(n + m)m$ while the number of equations in (9) is $(n + 1)m + (nm - |\Omega|)$. Elementary calculations lead to the bound n_0 .

The proof of the other two necessary conditions is based on the fact that

$$(H_0)_{i,:} Y = (X_0)_{i,:},$$

where $H_0 = D_0^{-1}$. To simplify the notations, we omit the subscript 0 from H_0 , D_0 and X_0 in the rest of this proof. Divide the sample matrix Y into two sub-matrices $Y_{:, \Omega_i}$ and $Y_{:, \Omega_i^c}$. Then it holds that

$$H_{i,:} Y_{:, \Omega_i} = 0^T, \text{ and } H_{i,:} Y_{:, j} \neq 0 \quad \forall j \in \Omega_i.$$

($H_{i,:}^T$ is in the null space of $Y_{:, \Omega_i^c}$.) Hence $H_{i,:}$ is unique (up to a scaling factor) if and only if $\text{span}(Y_{:, \Omega_i^c})$ has dimension $m - 1$. In this case, $\text{span}(H_{i,:}^T)$ is the null space of both $\text{span}(Y_{:, \Omega_i^c})$ and $\text{span}(D_{:, \{i\}^c})$. It is concluded that $H_{i,:}$ is unique if and only if $\text{span}(Y_{:, \Omega_i^c}) = \text{span}(D_{:, \{i\}^c})$.

Necessary condition 2 follows directly from that $\text{rank}(Y_{:, \Omega_i^c}) = m - 1$.

To prove the last necessary condition, note first that the fact that $\text{span}(Y_{:, \Omega_i^c}) = \text{span}(D_{:, \{i\}^c})$ implies that each column of $D_{:, \{i\}^c}$ participates in generating some columns of $Y_{:, \Omega_i^c}$. That is, $\forall i' \neq i$, $D_{:, i'}$ participates in generating $Y_{:, j}$ for some $j \notin \Omega_i$. Necessary condition 3 is therefore proved. Note that condition 3 is not sufficient. It does not prevent the following rank deficient case: there exist $i'_1, i'_2 \in \{i\}^c$ such that both $D_{:, i'_1}$ and $D_{:, i'_2}$ only participate in generating a single sample in $Y_{:, j}$ for some $j \in \Omega_i^c$. $\square$

3) *Discussions on the Number of Samples :* We now study the number of samples n needed to ensure that the necessary conditions for unique recovery, as specified in Proposition 2, hold with high probability. To that end we use the following probabilistic model: entries of $D_0 \in \mathbb{R}^{m \times m}$ are independently generated from the Gaussian distribution $\mathcal{N}(0, \frac{1}{m})$, and entries of $X_0 \in \mathbb{R}^{m \times n}$ are independently generated from the Bernoulli-Gaussian distribution $BG(\theta)$ with $\theta \in [0, 1]$, where Bernoulli-Gaussian distribution is defined as follows.

Definition 2. A random variable X is Bernoulli-Gaussian distributed $X \sim BG(\theta)$ with $\theta \in [0, 1]$, if $X = W \cdot C$, where random variables W and C are independent, W is Bernoulli distributed with parameter θ , and $C \sim \mathcal{N}(0, 1)$.

Remark 1. The Gaussian distribution is not essential. It can be replaced by any continuous distribution.

Proposition 3 (Number of Samples). Suppose that $Y = D_0 X_0$ where D_0 and X_0 are generated according to the above probability model. Given a constant $\epsilon \in (0, 1)$, the i -th necessary condition in Proposition 2 holds with probability at least $1 - \epsilon$, if $n \geq n_i$, where

$$n_1 = \frac{m-1}{1-\theta} \left[1 - \frac{\ln \epsilon}{4m(m-1)(1-\theta)} + \sqrt{\left(1 - \frac{\ln \epsilon}{4m(m-1)(1-\theta)} \right)^2 - 1} \right].$$

$$n_2 = \frac{m-1}{1-\theta} \left[1 - \frac{\ln \epsilon - \ln m}{4(m-1)(1-\theta)} + \sqrt{\left(1 - \frac{\ln \epsilon - \ln m}{4(m-1)(1-\theta)} \right)^2 - 1} \right].$$

and

$$n_3 = \frac{\ln \epsilon - \ln m - \ln(m-1)}{\ln(1-\theta(1-\theta))}.$$

Furthermore, it holds that $n_1 \leq n_2$. If $n \geq \max(n_2, n_3)$, then all three necessary conditions in Proposition 2 hold.

Proof. See Appendix A. $\square$

Remark 2. We have the following observations.

- With fixed ϵ and θ , n_1 and n_2 scale linearly with m while n_3 is proportional to $\ln m$.
- With fixed m and θ , n_1 , n_2 , and n_3 increase proportionally to $-\ln \epsilon$.
- With fixed m and ϵ , when θ increases from 0 to 1, n_1 and n_2 increase, while n_3 first decreases and then increases. This matches the intuition that when θ is too small, we need more samples to have enough information to recover the dictionary. On the other hand, when θ is too large, more samples are needed to generate the orthogonal space of each $H_{i,:}^T$. This is verified by simulations in Section V.

The bound $\max(n_2, n_3)$ provides a good estimate of the number of samples needed for unique recovery. By set theory, if event A is a necessary condition for B , then B implies A , or equivalently, $B \subseteq A$ and $\Pr(B) \leq \Pr(A)$. In Proposition 3, the quantity $1 - \epsilon$ is a lower bound for $\Pr(A)$, where these necessary conditions hold. But unfortunately it is neither lower nor upper bound for $\Pr(B)$, where the dictionary can be uniquely recovered. Nevertheless, our simulations show that $\max(n_2, n_3)$ is a good approximation to the number of samples needed to recover the dictionary uniquely with probability more than $1 - \epsilon$.

In an asymptotic regime, the bounds can be simplified.

Corollary 1 (Asymptotic Bounds). Consider the same settings as in Proposition 3. For a given $\theta \in (0, 1)$, let $m, n \rightarrow \infty$ with $\frac{n}{m} \rightarrow \bar{n} \in \mathbb{R}^+$. If $\bar{n} > \frac{1}{1-\theta}$, then all three necessary conditions in Proposition 2 holds with a probability arbitrary close to 1.

This corollary follows from elementary calculations and the fact that $\ln(m)/m \rightarrow 0$ when $m \rightarrow \infty$.

IV. DICTIONARY UPDATE WITH UNCERTAINTY

While Section III studies the ideal case, this section focuses on the general case using the insight from Section III. In practice, there may be noise in the training samples $\mathbf{Y}$, and there may be errors in the assumed sparsity pattern. The exact equality in (6) may not hold any longer. Following the idea in Section III, total least squares methods are applied to handle the uncertainties. The techniques for non-overcomplete and overcomplete dictionaries are developed in Sections IV-A and IV-B respectively.

A. Non-overcomplete Dictionary Update

In the case $m \geq l$, let $\mathbf{H} = \mathbf{D}^\dagger$ be the pseudo-inverse of $\mathbf{D}$ and assume that $\mathbf{H}\mathbf{D} = \mathbf{I}_{l \times l}$. Due to the uncertainty, Equation (9) becomes approximate, that is,

$$\mathbf{H}[\mathbf{Y}, \mathbf{1}] \approx [\mathbf{X}, \mathbf{1}] \text{ and } \mathcal{P}_{\Omega^c}(\mathbf{X}) \approx \mathbf{0}. \quad (10)$$

Total least squares is a technique to solve a least squares problem in the form $\mathbf{A}\mathbf{X} \approx \mathbf{B}$ where errors in both observations $\mathbf{B}$ and regression models $\mathbf{A}$ are considered [44], [45]. It targets at minimising the total errors via

$$\min_{\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \mathbf{X}} \|\mathbf{A} - \tilde{\mathbf{A}}, \mathbf{B} - \tilde{\mathbf{B}}\|_F^2, \text{ subject to } \tilde{\mathbf{A}}\mathbf{X} = \tilde{\mathbf{B}}. \quad (11)$$

The constraint set above is nonconvex and hence (11) is a nonconvex optimisation problem. Nevertheless, its global optimal solution can be obtained by using the singular value decomposition (SVD). Set $\mathbf{Z} = [\mathbf{X}^T, -\mathbf{I}]^T$. Observe that the constraint in (11) implies that $[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}]\mathbf{Z} = \mathbf{0}$. The optimal $\mathbf{Z}$ can be obtained from the smallest right singular vectors of the matrix $[\mathbf{A}, \mathbf{B}]$, and the optimal $[\tilde{\mathbf{A}}, \tilde{\mathbf{B}}]$ is a best lower-rank approximation of the matrix $[\mathbf{A}, \mathbf{B}]$.

The difficulty in applying total least squares directly is due to the additional constraint $\mathcal{P}_{\Omega^c}(\mathbf{X}) \approx \mathbf{0}$ in (10). Below we present three possible solutions, where the last one IterTLS excels and is adopted.

1) *Structured Total Least Squares (STLS)*: Consider having uncertainties in both $\mathbf{Y}$ and the sparsity pattern. Based on (10), a straightforward total least squares formulation is

$$\min_{\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H}} \frac{1}{2} \|\mathbf{Y} - \tilde{\mathbf{Y}}\|_F^2 + \frac{1}{2} \|\mathcal{P}_{\Omega^c}(\tilde{\mathbf{X}})\|_2^2, \quad (12)$$

s.t. $\mathbf{H}[\tilde{\mathbf{Y}}, \mathbf{1}] = [\tilde{\mathbf{X}}, \mathbf{1}]$.

To solve the above nonconvex optimisation problem, we follow the approach in [46]. It involves an iterative process where each iteration solves an approximated quadratic optimisation problem which admits a closed-form optimal solution.

At each iteration, denote the initial estimate of $(\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H})$ by $(\hat{\mathbf{Y}}, \hat{\mathbf{X}}, \hat{\mathbf{H}})$. Note that the constraint set in (12) can be written as $\mathcal{L}(\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H}) = \mathbf{0}$ where

$$\mathcal{L}(\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H}) := \mathbf{H}[\tilde{\mathbf{Y}}, \mathbf{1}] - [\tilde{\mathbf{X}}, \mathbf{1}].$$

We consider the first order Taylor approximation of $\mathcal{L}(\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H})$ at given $(\hat{\mathbf{Y}}, \hat{\mathbf{X}}, \hat{\mathbf{H}})$, which reads

$$\mathcal{L}(\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H}) \approx \mathcal{L}(\hat{\mathbf{Y}}, \hat{\mathbf{X}}, \hat{\mathbf{H}}) + \mathbf{J}(\mathbf{z} - \hat{\mathbf{z}}),$$

where $\mathbf{z} := [\text{vect}(\tilde{\mathbf{Y}})^T, \text{vect}(\tilde{\mathbf{X}})^T, \text{vect}(\mathbf{H})^T]^T$, $\hat{\mathbf{z}} := [\text{vect}(\hat{\mathbf{Y}})^T, \text{vect}(\hat{\mathbf{X}})^T, \text{vect}(\hat{\mathbf{H}})^T]^T$, and $\mathbf{J}$ is the corresponding Jacobian matrix. With this approximation, the nonconvex optimisation problem in (12) becomes a quadratic optimisation problem with equality constraints

$$\min_{\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H}} \frac{1}{2} \|\mathbf{Y} - \tilde{\mathbf{Y}}\|_F^2 + \frac{1}{2} \|\mathcal{P}_{\Omega^c}(\tilde{\mathbf{X}})\|_2^2, \quad (13)$$

s.t. $\mathcal{L}(\hat{\mathbf{Y}}, \hat{\mathbf{X}}, \hat{\mathbf{H}}) + \mathbf{J}(\mathbf{z} - \hat{\mathbf{z}}) = \mathbf{0}$.

This is a quadratic optimisation problem with linear equality constraints, and admits a closed-form solution by a direct application of KKT conditions [47].

The STLS approach has two issues. The first issue is its very high computational cost. The quadratic optimisation problem (13) involves $mn + ln + lm$ unknowns and $l(n + 1)$ equation constraints. Its closed-form solution involves a matrix of size $(n(m + 2l) + l(m + 1)) \times (n(m + 2l) + l(m + 1))$. We have obtained the closed-form of the Jacobian matrix $\mathbf{J}$, implemented a conjugate gradient algorithm to use the structures in (13) for a speed-up (details are omitted here). However, simulations in Section V-B show that the computation speed is still too slow for practical problems. The second issue is the inferior performance compared to other TLS methods in Sections IV-A2 and IV-A3. This is because Taylor approximation of the constraint is used in STLS, while other TLS methods below incorporate the constraints directly without Taylor approximation.

2) *Parallel Total Least Squares (ParTLS)*: The key idea of ParTLS is to decouple the problem (10) into l sub-problems that can be solved in parallel:

$$\mathbf{H}_{i,:}[\mathbf{Y}, \mathbf{1}] \approx [\mathbf{X}_{i,:}, \mathbf{1}] \text{ and } \mathbf{X}_{i,\Omega_i^c} \approx \mathbf{0}, \forall i \in [l].$$

It is straightforward to verify that this is equivalent to

$$\underbrace{\begin{bmatrix} \mathbf{Y}^T & -\mathbf{P}_{\Omega_i} \\ \mathbf{1} & \mathbf{0} \end{bmatrix}}_{\mathbf{A}_i} \underbrace{\begin{bmatrix} \mathbf{H}_{i,:}^T \\ \mathbf{X}_{i,\Omega_i}^T \end{bmatrix}}_{\mathbf{z}_i} \approx \underbrace{\begin{bmatrix} \mathbf{0} \\ \mathbf{1} \end{bmatrix}}_{\mathbf{b}}, \forall i \in [l], \quad (14)$$

where $\mathbf{P}_{\Omega_i} \in \mathbb{R}^{n \times |\Omega_i|}$ is the projection matrix obtained by keeping the columns of the identity matrix indexed by Ω_i and removing all other columns.

Sub-problems (14) can be solved by directly applying the TLS formulation (11). Note that $[\tilde{\mathbf{A}}_i, -\tilde{\mathbf{b}}][\mathbf{z}_i^T, \mathbf{1}]^T = \mathbf{0}$. The vector $[\mathbf{z}_i^T, \mathbf{1}]^T$ can be computed as a scaled version of the least right singular vector of the matrix $[\mathbf{A}_i, -\mathbf{b}]$. Then $\mathbf{H}_{i,:}$ and $\mathbf{X}_{i,\Omega_i}$ can be obtained from $\mathbf{z}_i$.

ParTLS enjoys the following advantages. 1) Its global optimality is guaranteed for the ideal case of no data noise or sparsity pattern errors. It is straightforward to see that in the ideal case the ParTLS solutions satisfy (9). 2) It is computationally efficient. The sub-problems (14) are of small size and can be solved in parallel. However, ParTLS also has its own issue — certain structures in the problem are not

enforced. For different sub-problem $i \in [L]$, the ‘denoised’ $\mathbf{Y}$, denoted by $\tilde{\mathbf{Y}}_i$ can be different.

3) *Iterative Total Least Squares (IterTLS)* : IterTLS is an iterative algorithm such that in each iteration a total least squares problem is formulated based on the estimate from the previous iteration. It starts with an initial estimate obtained by solving the ideal case equation (9). In each iteration, let $\hat{\mathbf{X}}$ be an estimate of $\mathbf{X}$ from either initialisation or the previous iteration. We formulate the following total least squares problem

$$\min_{\tilde{\mathbf{Y}}, \tilde{\mathbf{X}}, \mathbf{H}} \left\| [\mathbf{Y}^T - \tilde{\mathbf{Y}}^T, \hat{\mathbf{X}}^T - \tilde{\mathbf{X}}^T] \right\|_F^2 \text{ s.t. } \tilde{\mathbf{Y}}^T \mathbf{H}^T = \tilde{\mathbf{X}}^T, \quad (15)$$

which has the identical form as (11). Note that the constraint $\mathcal{P}_{\Omega^c}(\mathbf{X}) \approx \mathbf{0}$ in (10) is implicitly imposed as $\mathcal{P}_{\Omega^c}(\hat{\mathbf{X}}) = \mathbf{0}$. The problem (15) can be optimally solved by using the SVD

$$[\mathbf{Y}^T, \hat{\mathbf{X}}^T] = [\mathbf{U}_Y, \mathbf{U}_X] \begin{bmatrix} \Sigma_Y & \mathbf{0} \\ \mathbf{0} & \Sigma_X \end{bmatrix} \begin{bmatrix} \mathbf{V}_{YY} & \mathbf{V}_{YX} \\ \mathbf{V}_{XY} & \mathbf{V}_{XX} \end{bmatrix}^T.$$

The optimal solution is given by $\tilde{\mathbf{Y}}^T = \mathbf{U}_Y \Sigma_Y \mathbf{V}_{YY}^T$ and $\tilde{\mathbf{X}}^T = \mathbf{U}_Y \Sigma_Y \mathbf{V}_{XY}^T$. To prepare the next iteration, one obtains an updated estimate $\hat{\mathbf{X}}'$ by applying a simple projection operator to $\tilde{\mathbf{X}}$: $\mathcal{P}_{\Omega}(\hat{\mathbf{X}}') = \mathcal{P}_{\Omega}(\tilde{\mathbf{X}})$ and $\mathcal{P}_{\Omega^c}(\hat{\mathbf{X}}') = \mathbf{0}$. With this new estimate $\hat{\mathbf{X}}'$, one can proceed with the next iteration until convergence.

B. Update Overcomplete Dictionary

The difficulty of overcomplete dictionary update comes from the fact that for an overcomplete dictionary $\mathbf{D}_0$, $\mathbf{D}_0^\dagger \mathbf{D}_0 \neq \mathbf{I}$ in general, where $\mathbf{D}_0^\dagger$ is the pseudo-inverse of $\mathbf{D}_0$. Therefore, the above least squares or total least squares approaches cannot be directly applied.

To address this issue, a straightforward approach is to divide the whole dictionary into a set of sub-dictionaries each of which is either complete or undercomplete, and then update these sub-dictionaries one-by-one whilst fixing all other sub-dictionaries and the corresponding coefficients. More explicitly, given estimated $\mathbf{D}$ and $\mathbf{X}$, consider updating $\mathbf{D}_{:, \mathcal{T}}$, the submatrix of $\mathbf{D}$ consisting of columns indexed by $\mathcal{T}$, and $\mathbf{X}_{\mathcal{T}, :}$, the submatrix of $\mathbf{X}$ consisting of rows indexed by $\mathcal{T}$. Then, consider the residual matrix

$$\mathbf{Y}_r = \mathbf{Y} - \mathbf{D}_{:, \mathcal{T}^c} \mathbf{X}_{\mathcal{T}^c, :}, \quad (16)$$

and apply the method in Section IV-A3 to solve the problem $\mathbf{Y}_r \approx \mathbf{D}_{:, \mathcal{T}} \mathbf{X}_{\mathcal{T}, :}$. Then repeat this step for all sub-dictionaries. As the dictionary is updated block by block, we refer to our algorithm as BLock Total LEast SquareS (BLOTLESS).

V. NUMERICAL TEST

Parts of the numerical tests are based on synthetic data. The training samples $\mathbf{Y} = \mathbf{D}_0 \mathbf{X}_0$ are generated according to the probability model specified in Section III-3. When the dictionary recovery is not exact, the performance criterion is the difference between the ground-truth dictionary $\mathbf{D}_0$ and the

estimated dictionary $\hat{\mathbf{D}}$. In particular, the estimation error is defined as

$$R_{\text{err}} = \frac{1}{l} \sum_{p=1}^l \left(1 - |\hat{\mathbf{d}}_p^T \mathbf{d}_{0_{j_p}}| \right), \quad (17)$$

where $\hat{\mathbf{d}}_p$ is the p -th item in the estimated dictionary, $\mathbf{d}_{0_{j_p}}$ is the j_p -th item in the ground-truth dictionary, $j_p = \arg \max_{j \in \mathcal{J}_p} \{ |\hat{\mathbf{d}}_p^T \mathbf{d}_{0_j}| \}$ and $\mathcal{J}_p = [L] \setminus \{j_1, j_2, \dots, j_{p-1}\}$. The items in both dictionaries are normalised to have unit ℓ_2 -norm.

Numerical tests based on real data are presented in Section V-D2 for image denoising. The performance metric is the Peak Signal-to-Noise-Ratio (PSNR) of the denoised images.

A. Simulations for Exact Dictionary Recovery

In this section we evaluate numerically bounds in Proposition 3. All the results presented here are based on 100 random and independent trials. For theoretical performance prediction, we compute n_2 , n_3 , and $\max(n_2, n_3)$ using $\epsilon = 0.01$. In the numerical simulations, we vary n and find its critical value n_{sim} under which exact recovery happens with an empirical probability at most 99% and above which exact recovery happens with an empirical probability at least 99%.

Fig. 1: Probability of exact recovery against the number of training samples for $m = 30$.

We start with the relation between the number of training samples n and the sparsity ratio θ for a given m . In Figure 1, we fix $m = 30$, vary θ , and study the probability of exact recovery against the number of training samples. Results in Figure 1 show that the theoretical prediction $\max(n_2, n_3)$ is quite close to the critical n_{sim} obtained by simulations. One can also observe that the needed number of training samples first decreases and then increases as θ increases, which is also predicted by the theoretical bounds. A larger scale numerical test is done in Figure 2, where four sub-figures correspond to four different values of m . Once again,

Fig. 2: Normalised number of training samples n/m for at least 99% probability of exact recovery versus m .

Fig. 3: Required n for exact recovery versus m with a given $\theta = 0.2$.

simulations demonstrate these bounds in Proposition 3 match the simulations very well.

Let us consider the required n for exact recovery as a function of m by fixing $\theta = 0.2$. Simulation results are depicted in Figure 3. When $n_3 \geq n_2$, n_{sim} behaves as $\ln m$. Otherwise, n_{sim} behaves as m . This is consistent with Proposition 3.

Finally, we numerically study the accuracy of the asymptotic results in Corollary 1. We draw normalised number of training samples n/m for exact recovery as a function of sparsity ratio θ for different values of m , including $m = 16, 32, 64$. Simulation results in Figure 4 show a trend that is consistent with the asymptotic results in Corollary 1.

B. Total Least Squares Methods

The three total least squares methods introduced in Section IV are compared, henceforth denoted by BLOTLESS-STLS, BLOTLESS-ParTLS and BLOTLESS-IterTLS respectively, by being embedded in the dictionary learning process. Random

Fig. 4: Asymptotic bounds in Corollary 1 compared with empirical results for finite m .

dictionaries are used as the initial starting point of dictionary learning. OMP is used for sparse coding.

(a) $m = l = 64, \theta = 5/64, n = 400$ (b) $m = 64, l = 128, \theta = 5/128, n = 800$

Fig. 5: Performance comparison of different total least squares methods when used in the overall dictionary learning process. Results are averages of 50 independent trials.

TABLE I: Runtime (seconds per iteration of the dictionary learning process) comparison for different total least squares methods: $m = l = 64, \theta = 5/64$, and number of iterations $n_{it} = 50$. Results are averages of 50 independent trials using Matlab 2018b on a MacBook Pro with 16GB RAM and 2.3 GHz Intel Core i5 processor.

	$n = 200$	$n = 300$	$n = 400$	$n = 500$
STLS	621.6544	836.3027	1098.5955	1265.1895
ParTLS	9.2544	18.1191	25.3678	31.0551
IterTLS	6.3446	8.5157	10.9026	13.8056

Fig. 5 compares the dictionary learning errors (17) for both complete and over-complete dictionaries. BLOTLESS-IterTLS converges the fastest and has the smallest error floor. Then a runtime comparison is given in Table I. BLOTLESS-IterTLS clearly outperforms the other two methods. It is therefore used as the default dictionary update method in later simulations.

C. Robustness to Errors in Sparsity Pattern

Simulations are next designed to evaluate the robustness of different dictionary update algorithms to errors in sparsity pattern. Towards that end, a fraction r of indices in the true support are randomly chosen to be replaced with the same number of randomly chosen indices not in the true support

Fig. 6: Performance comparison of different update algorithm when different sparse pattern mismatch is applied. $m = l = 64$, $\theta = 5/64$, $n = 500$. Results are averages of 100 independent trials.

set. This erroneous sparsity pattern is then fed into different dictionary update algorithms. The numerical results in Fig. 6 demonstrate the robustness of BLOTLESS (with IterTLS).

D. Dictionary Learning with BLOTLESS Update

This subsection compares dictionary learning performance for different dictionary update methods. The sparse coding algorithm is OMP. IterTLS in Section IV-A3 is used for BLOTLESS. Results for synthetic data are presented in Section V-D1 while Section V-D2 focuses on image denoising using real data.

(a) $m = l = 64$, $n = 400$, $\theta = 5/64$. (b) $m = 64$, $l = 128$, $n = 600$, $\theta = 5/128$.

(c) $m = l = 64$, $\theta = 5/64$, $n_{it} = 150$. (d) $m = 64$, $l = 128$, $\theta = 5/128$, $n_{it} = 150$.

Fig. 7: Comparison of dictionary update methods for the noise free case. Results are averages of 100 trials.

1) *Synthetic Data* : Fig. 7 and 8 compare the performance of dictionary learning using different dictionary update algorithms. Fig. 7 focus on the noise free case where $\mathbf{Y} = \mathbf{D}_0\mathbf{X}_0$ and Fig. 8 concerns with the noisy case where

Fig. 8: Comparison of dictionary update methods for the noisy case: SNR is 15dB. Results are averages of 100 trials. n_{it} denotes the number of iterations.

$\mathbf{Y} = \mathbf{D}_0\mathbf{X}_0 + \mathbf{W}$ where $\mathbf{W}$ is the additive Gaussian noise matrix with i.i.d. entries and the signal-to-noise ratio (SNR) is set to 15dB. Both figures include the cases of complete and over-complete dictionaries. The results presented in Fig. 7 and 8 clearly indicate that dictionary learning based on BLOTLESS converges much faster and needs at least 1/3 less training samples than other benchmark dictionary update methods.

Fig. 9: Dictionary learning via BLOTLESS with different block sizes: $m = l = 64$, $n = 200$, $\theta = 5/64$.

In BLOTLESS update, blocks of the dictionary (sub-dictionaries) are updated sequentially. Fig. 9 compares the performance with different block sizes. Note that when each block contains only one dictionary item, the dictionary update problem is the same as that in K-SVD. Hence the performance of K-SVD is added in Fig. 9. Simulations suggest that the larger the dictionary blocks are, the faster the convergence is and the better performance is. The performance of BLOTLESS

with block size one is slightly better than that of K-SVD. This is because the TLS step in IterTLS does not enforce the sparsity pattern and hence better accommodates errors in the estimated sparsity pattern.

2) *Real Data* : We use the Olivetti Research Laboratory (ORL) face database [48] for dictionary learning and then use the learned dictionary for image denoising.

For dictionary learning, according to the simulation results in Section V-D1, $n = 500$ samples of size 8×8 patches from face images are enough for training a dictionary via BLOTLESS. The parameters used in dictionary learning are $m = 64$, $l = 128$, $\theta = 10/128$, and $n_{it} = 20$. After learning a dictionary, image denoising [1] is performed using test images from the same dataset. The denoising results are shown in Table II, where four test images are used. In all these four tests, the BLOTLESS method outperforms all other algorithms, which is consistent with these simulations in V-D1.

VI. CONCLUSION

This paper proposed a BLOTLESS algorithm for dictionary update. It divides the dictionary into sub-dictionaries, each of which is non-overcomplete. Then BLOTLESS updates a sub-dictionary and the corresponding sparse coefficients using least squares or total least squares approaches. Necessary conditions for unique recovery are identified and they hold with high probability when the number of training samples is larger than the derived bounds in Proposition 3. Simulations show that these bounds match the simulations well, and that BLOTLESS outperforms other benchmark algorithms. One future direction is to find sufficient bounds for unique recovery and their comparisons to the necessary bounds.

APPENDIX

A. Proof of Proposition 3

The proof needs Hoeffding's inequality [49] for Bernoulli random variables, stated below.

Lemma 1 (Hoeffding's Inequality). *For N many identical Bernoulli random variables $\{X_i, i = 1, 2, \dots, N\}$. Each X_i takes the value 1 with probability p and 0 with probability $(1 - p)$, then the following Hoeffding's inequality holds*

$$\Pr \left(\sum_{i=1}^N X_i \geq (p + \lambda)N \right) \leq \exp(-2\lambda^2 N), \quad (18)$$

where $\lambda > 0$ is a constant number.

To derive n_1 , we consider the case that the necessary condition 1 in Proposition 2 fails. That is,

$$\begin{aligned} |\Omega| &\geq nm - m^2 + m \\ &= nm\left(\theta + \left(1 - \theta - \frac{m}{n} + \frac{1}{n}\right)\right). \end{aligned}$$

Based on Hoeffding's inequality, the probability of this event is upper bounded by

$$\exp \left(-2 \left(1 - \theta - \frac{m}{n} + \frac{1}{n} \right)^2 mn \right).$$

If this probability is smaller than ϵ , it follows that

$$((1 - \theta)n - (m - 1))^2 + \frac{\ln \epsilon}{2m}n \geq 0.$$

The left hand side of the above inequality is quadratic in n , which after some elementary calculations leads to

$$\begin{aligned} n \geq n_1 &= \frac{m - 1}{1 - \theta} \left[1 - \frac{\ln \epsilon}{4m(m - 1)(1 - \theta)} \right. \\ &\quad \left. + \sqrt{\left(1 - \frac{\ln \epsilon}{4m(m - 1)(1 - \theta)} \right)^2 - 1} \right]. \end{aligned}$$

The derivation of n_2 is similar. Consider the probability that the necessary condition 2 in Proposition 2 fails:

$$\begin{aligned} &1 - \Pr(\forall i \in [m], |\Omega_i^c| \geq m - 1) \\ &= \Pr(\exists i \in [m], |\Omega_i| \geq n - m + 1) \\ &\leq m \Pr(|\Omega_1| \geq n - m + 1) \\ &= m \Pr \left(|\Omega_1| \geq n \left(\theta + \left(1 - \theta - \frac{m - 1}{n} \right) \right) \right), \end{aligned}$$

where the inequality in the third line follows from the union bound. After applying Hoeffding's inequality and setting the upper bound less than ϵ we obtain

$$((1 - \theta)n - (m - 1))^2 + \frac{\ln \epsilon - \ln m}{2}n \geq 0.$$

It follows that

$$\begin{aligned} n \geq n_2 &= \frac{m - 1}{1 - \theta} \left[1 - \frac{\ln \epsilon - \ln m}{4(m - 1)(1 - \theta)} \right. \\ &\quad \left. + \sqrt{\left(1 - \frac{\ln \epsilon - \ln m}{4(m - 1)(1 - \theta)} \right)^2 - 1} \right]. \end{aligned}$$

To derive n_3 , we define the following event

- $S_{i,i'}$: For given $i \neq i'$, $\exists j \in [n]$ such that $X_{i,j} = 0$ and $X_{i',j} \neq 0$.

Then the probability that the necessary condition 3 fails is given by

$$\begin{aligned} 1 - \Pr(\cap_{i \neq i'} S_{i,i'}) &= \Pr(\cup_{i \neq i'} S_{i,i'}^c) \\ &\leq m(m - 1) \Pr(S_{1,2}^c) \\ &= m(m - 1)(1 - \theta(1 - \theta))^n, \end{aligned}$$

where the inequality in the second line follows from the union bound. If we set this probability to be smaller than ϵ , we obtain

$$n \geq n_3 = \frac{\ln \epsilon - \ln m - \ln(m - 1)}{\ln(1 - \theta(1 - \theta))}.$$

REFERENCES

- [1] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Trans. Image Process.*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [2] L. Liu, L. Chen, C. P. Chen, Y. Y. Tang *et al.*, "Weighted joint sparse representation for removing mixed noise in image," *IEEE Trans. Cybern.*, vol. 47, no. 3, pp. 600–611, 2017.
- [3] M. M. Bronstein, A. M. Bronstein, M. Zibulevsky, and Y. Y. Zeevi, "Blind deconvolution of images using optimal sparse representations," *IEEE Trans. Image Process.*, vol. 14, no. 6, pp. 726–736, 2005.

TABLE II: Denoising comparison using different learnt dictionaries, where the denoised PSNR (dB) are computed and shown in table.

Original Image
Noisy Image	28.13 dB	22.11 dB	18.59 dB	28.13 dB	22.11 dB	18.59 dB	28.13 dB	22.11 dB	18.59 dB	28.13 dB	22.11 dB	18.59 dB
MOD	33.00 dB	29.24 dB	27.22 dB	32.68 dB	28.84 dB	26.82 dB	31.95 dB	27.43 dB	25.76 dB	32.13 dB	28.23 dB	26.26 dB
K-SVD	32.50 dB	28.72 dB	26.79 dB	32.03 dB	28.26 dB	26.35 dB	31.49 dB	27.43 dB	25.31 dB	31.58 dB	27.74 dB	25.81 dB
SimCO	33.43 dB	29.78 dB	27.81 dB	33.58 dB	30.11 dB	28.04 dB	32.18 dB	28.51 dB	26.47 dB	32.65 dB	29.18 dB	27.27 dB
BLOTLESS	33.67 dB	29.90 dB	27.95 dB	33.91 dB	30.33 dB	28.20 dB	32.38 dB	28.67 dB	26.66 dB	32.88 dB	29.38 dB	27.45 dB

- [4] J. Yang, J. Wright, T. S. Huang, and Y. Ma, "Image super-resolution via sparse representation," *IEEE Trans. Image Process.*, vol. 19, no. 11, pp. 2861–2873, 2010.
- [5] Q. Dai, S. Yoo, A. Kappeler, and A. K. Katsaggelos, "Sparse representation-based multiple frame video super-resolution," *IEEE Trans. Image Process.*, vol. 26, no. 2, pp. 765–781, 2017.
- [6] Z. Cvetkovic, "On discrete short-time fourier analysis," *IEEE Trans. Signal Process.*, vol. 48, no. 9, pp. 2628–2640, 2000.
- [7] Z. Cvetkovic and M. Vetterli, "Discrete-time wavelet extrema representation: design and consistent reconstruction," *IEEE Trans. Signal Process.*, vol. 43, no. 3, pp. 681–693, 1995.
- [8] E. J. Candes and D. L. Donoho, "Curvelets: A surprisingly effective nonadaptive representation for objects with edges," Stanford Univ Ca Dept of Statistics, Tech. Rep., 2000.
- [9] B. A. Olshausen and D. J. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, p. 607, 1996.
- [10] M. Aharon, M. Elad, A. Bruckstein *et al.*, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, p. 4311, 2006.
- [11] K. Engan, S. O. Aase, and J. H. Husoy, "Method of optimal directions for frame design," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, vol. 5, 1999, pp. 2443–2446.
- [12] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski, "Dictionary learning algorithms for sparse representation," *Neur. Comput.*, vol. 15, no. 2, pp. 349–396, 2003.
- [13] W. Dai, T. Xu, and W. Wang, "Simultaneous codeword optimization (SimCO) for dictionary update and learning," *IEEE Trans. Signal Process.*, vol. 60, no. 12, pp. 6340–6353, 2012.
- [14] I. Tomic and P. Frossard, "Dictionary learning," *IEEE Signal Process. Mag.*, vol. 28, no. 2, pp. 27–38, 2011.
- [15] J. A. Tropp and A. C. Gilbert, "Signal recovery from random measurements via orthogonal matching pursuit," *IEEE Trans. Inf. Theory*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [16] D. Needell and R. Vershynin, "Uniform uncertainty principle and signal recovery via regularized orthogonal matching pursuit," *Found. of Comput. Math.*, vol. 9, no. 3, pp. 317–334, 2009.
- [17] W. Dai and O. Milenkovic, "Subspace pursuit for compressive sensing signal reconstruction," *IEEE Trans. Inf. Theory*, vol. 55, no. 5, pp. 2230–2249, 2009.
- [18] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Rev.*, vol. 43, no. 1, pp. 129–159, 2001.
- [19] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [20] J. A. Tropp and S. J. Wright, "Computational methods for sparse solution of linear inverse problems," *Proc. IEEE*, vol. 98, no. 6, pp. 948–958, 2010.
- [21] S. O. Aase, J. H. Husoy, J. Skretting, and K. Engan, "Optimized signal expansions for sparse representation," *IEEE Trans. Signal Process.*, vol. 49, no. 5, pp. 1087–1096, 2001.
- [22] K. Skretting, K. Engan, J. Husoy, and S. O. Aase, "Sparse representation of images using overlapping frames," in *Proc. of the scandinavian Conf. image anal.*, 2001, pp. 613–620.
- [23] K. Engan, K. Skretting, and J. H. Husoy, "Family of iterative ls-based dictionary learning algorithms, ILS-DLA, for sparse signal representation," *Digital Signal Process.*, vol. 17, no. 1, pp. 32–49, 2007.
- [24] L. N. Smith and M. Elad, "Improving dictionary learning: Multiple dictionary updates and coefficient reuse," *IEEE Signal Process. Lett.*, vol. 20, no. 1, pp. 79–82, 2013.
- [25] M. Aharon, M. Elad, and A. M. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear Algeb. and Its Appl.*, vol. 416, no. 1, pp. 48–67, 2006.
- [26] C. J. Hillar and F. T. Sommer, "When can dictionary learning uniquely recover sparse data from subsamples?" *IEEE Trans. Inf. Theory*, vol. 61, no. 11, pp. 6290–6297, 2015.
- [27] R. Remi and K. Schnass, "Dictionary identificationsparse matrix-factorization via l_1 -minimization," *IEEE Trans. Inf. Theory*, vol. 56, no. 7, pp. 3523–3539, 2010.
- [28] Q. Geng and J. Wright, "On the local correctness of l_1 -minimization

- for dictionary learning,” in *IEEE Int. Symp. Inf. Theory (ISIT)*, 2014, pp. 3180–3184.
- [29] K. Schnass, “Local identification of overcomplete dictionaries,” *J. Mach. Learn. Res.*, vol. 16, pp. 1211–1242, 2015.
- [30] K. Schnass, “On the identifiability of overcomplete dictionaries via the minimisation principle underlying K-SVD,” *Appl. and Comput. Harmon. Anal.*, vol. 37, no. 3, pp. 464–491, 2014.
- [31] K. Schnass, “Convergence radius and sample complexity of ITKM algorithms for dictionary learning,” *Appl. and Comput. Harmon. Anal.*, 2016.
- [32] D. A. Spielman, H. Wang, and J. Wright, “Exact recovery of sparsely-used dictionaries,” in *Conf. Learn. Theory*, 2012, pp. 37–1.
- [33] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, “Learning sparsely used overcomplete dictionaries,” in *Conf. Learn. Theory*, 2014, pp. 123–137.
- [34] A. Agarwal, A. Anandkumar, and P. Netrapalli, “Exact recovery of sparsely used overcomplete dictionaries,” *Stat.*, vol. 1050, pp. 8–39, 2013.
- [35] S. Arora, R. Ge, and A. Moitra, “New algorithms for learning incoherent and overcomplete dictionaries,” in *Conf. Learn. Theory*, 2014, pp. 779–806.
- [36] S. Arora, R. Ge, T. Ma, and A. Moitra, “Simple, efficient, and neural algorithms for sparse coding,” *Proc. Mach. Learn. Res.*, 2015.
- [37] B. Barak, J. A. Kelner, and D. Steurer, “Dictionary learning and tensor decomposition via the sum-of-squares method,” in *Proc. the 47th ann. ACM symp. Theory of comput.* ACM, 2015, pp. 143–151.
- [38] S. Arora, A. Bhaskara, R. Ge, and T. Ma, “More algorithms for provable dictionary learning,” *arXiv preprint arXiv:1401.0579*, 2014.
- [39] J. Sun, Q. Qu, and J. Wright, “Complete dictionary recovery over the sphere i: Overview and the geometric picture,” *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 853–884, 2017.
- [40] J. Sun, Q. Qu, and J. Wright, “Complete dictionary recovery over the sphere ii: Recovery by riemannian trust-region method,” *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 885–914, 2017.
- [41] S. Ling and T. Strohmer, “Self-calibration and bilinear inverse problems via linear least squares,” *SIAM J. Imag. Sci.*, vol. 11, no. 1, pp. 252–292, 2018.
- [42] R. Gribonval, G. Chardon, and L. Daudet, “Blind calibration for compressed sensing by convex optimization,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2012, pp. 2713–2716.
- [43] Q. Yu, W. Dai, Z. Cvetkovic, and J. Zhu, “Bilinear dictionary update via linear least squares,” in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2019, pp. 7923–7927.
- [44] I. Markovsky and S. Van Huffel, “Overview of total least-squares methods,” *Signal Process.*, vol. 87, no. 10, pp. 2283–2302, 2007.
- [45] S. Rhode, K. Usevich, I. Markovsky, and F. Gauterin, “A recursive restricted total least-squares algorithm,” *IEEE Trans. Signal Process.*, vol. 62, no. 21, pp. 5652–5662, 2014.
- [46] P. Lemmerling, N. Mastronardi, and S. Van Huffel, “Efficient implementation of a structured total least squares based speech compression method,” *Linear Algeb. and its Appl.*, vol. 366, pp. 295–315, 2003.
- [47] R. Fletcher, *Practical methods of optimization; 2nd ed.* Hoboken, NJ: Wiley, 2013.
- [48] F. Samaria and A. Harter, “Parameterisation of a stochastic model for human face identification,” in *Proc. IEEE Workshop Appl. Comp. Vision*, 1994, pp. 138–142.
- [49] N. I. Fisher and P. K. Sen, “Probability inequalities for sums of bounded random variables,” *J. of the Amer. Statist. Assoc.*, vol. 58, no. 301, pp. 13–30, 1963.