# Efficient and Parallel Separable Dictionary Learning

1st Cristian Rusu

Department of Automatic Control and Computers Faculty of Automatic Control and Computers, University Politehnica Bucharest Bucharest, Romania Research Center for Logic, Optimization and Security (LOS), Department of Computer Science Faculty of Mathematics and Computer Science, University of Bucharest Bucharest, Romania cristian.rusu@upb.ro 2<sup>nd</sup> Paul Irofti

Research Center for Logic, Optimization and Security (LOS), Department of Computer Science Faculty of Mathematics and Computer Science, University of Bucharest Bucharest, Romania paul@irofti.net, 0000-0002-7541-4334

Abstract—Separable, or Kronecker product, dictionaries provide natural decompositions for 2D signals, such as images. In this paper, we describe a highly parallelizable algorithm that learns such dictionaries which reaches sparse representations competitive with the previous state of the art dictionary learning algorithms from the literature but at a lower computational cost. We highlight the performance of the proposed method to sparsely represent image and hyperspectral data, and for image denoising.

*Index Terms*—sparse representations, dictionary learning, separable dictionaries, parallel computing, distributed computing

# I. INTRODUCTION

Dictionary learning (DL) [1] aims at finding a suitable overcomplete basis, or dictionary, that best represents a given dataset: training signals vectorized as columns in Y. Based on this training matrix we find the dictionary D that produces representations X such that  $Y \approx DX$ . We mandate sparsity in the representations [2] which implies that a sample y, uses only a few columns, also called atoms, from D encoded in the representations. Popular algorithms include Orthogonal Matching Pursuit (OMP) [3] for sparse representations and Approximate K-SVD (AK-SVD) [4] or more recent adaptive stable techniques [5] for dictionary learning.

Often, samples have an intrinsic structure that is lost or weakened with standard DL approaches. For example, in images and videos, pixels and voxels have strong vicinity based connections leading to certain patterns exploited by standard signal processing tools. Also, network or graph generated signals suffer from the same loss of structure and recent DL work [6], [7] has shown that recovering the graph structure through the learning process can significantly improve results. Thus, existing algorithms were adapted to separable variants that maintain the data structure by training separate dictionaries for each dimension [8]–[10] but also exploit the structure in order to gain better and faster approximations [11], [12].

Previous work introduced parallel/distributed DL algorithms but these addressed non-separable problems [13], [14] and introduced significant computational overhead [15]. We continue this pursuit and propose a new parallel algorithm that provides a fast and distributed (with little communication overhead) separable DL solution based on the Tensor MOD (TMOD) approach [16].

In the separable dictionary setting, for a bidimensional (2D) signal Y we now want to represent it using two dictionaries  $Y \approx D_1 X D_2^T$ . The approximation quality is identical  $||Y - D_1 X D_2^T||_F = ||\operatorname{vec} Y - (D_2 \otimes D_1)\operatorname{vec} X||_2$ , if the one-dimensional (1D) dictionary is set as  $D = D_2 \otimes D_1$ . Here  $\operatorname{vec} M$  denotes the column-order vectorization of matrix M,  $||.||_F$  the Frobenius norm and  $\otimes$  is the Kronecker product. Throughout the paper we treat the 2D case, but the results shown can be easily generalized to multiple dimensions [17].

In the context of DL, separable or pairwise dictionaries are used to represent 2D signals such as images. The data is now stored as a set of N samples  $Y_k \in \mathbb{R}^{m \times m}$ , k = 1 : N. Small dictionaries  $D_1 \in \mathbb{R}^{m \times n_1}$  (the left dictionary) and  $D_2 \in \mathbb{R}^{m \times n_2}$ (the right dictionary) cater to the 2D data such that they produce the corresponding  $X_k \in \mathbb{R}^{n_1 \times n_2}$  sparse representations that are  $s \in \mathbb{N}^*$  sparse (the matrix  $X_k$  has at most *s* nonzero elements). We are now ready to formulate the pairwise DL optimization problem as

$$\begin{array}{ll} \underset{D_{1}, D_{2}, X}{\text{minimize}} & \sum_{k=1}^{N} \|Y_{k} - D_{1}X_{k}D_{2}^{T}\|_{F}^{2} \\ \text{subject to} & \|X_{k}\|_{0} \leq s, \ k = 1:N, \\ & \|d_{1i}\|_{2} = 1 \text{ and } \|d_{2j}\|_{2} = 1, \end{array}$$
(1)

where  $i = 1 : n_1, j = 1 : n_2, d_{1i}$  and  $d_{2j}$  denote the normalized

C. Rusu was supported by a grant of the Ministry of Research, Innovation and Digitization, CNCS/CCCDI-UEFISCDI, project number PN-III-P1-1.1-TE-2019-1843, within PNCDI III. P. Irofti was supported by two grants of the Ministry of Research, Innovation and Digitization, CNCS/CCCDI-UEFISCDI, project number PN-III-P1-1.1-PD-2019-0825 and project number PN-III-P2-2.1-PED-2019-3248, within PNCDI III.

atoms of the two dictionaries. Existing algorithms [11], [12] exploit the separable formulation such that  $D_2 \otimes D_1$  is never explicitly computed.

### II. SEPARABLE DICTIONARY LEARNING

The earliest work on separable dictionaries, either in the context of dictionary learning or compressed sensing for images, was introduced in [18] and [19]. Recently, there has been revived interested in the separable dictionary learning problem, both theoretically and from an algorithmic perspective.

SeDiL [8] was among the first to attack the separable DL problem by employing a highly complex gradient descent algorithm on smooth Riemannian submanifolds. Keeping the data structure intact, thus multi-dimensional, has lead to many tensor-based algorithms [9], [16], [20], [21] that employ various CANDECOMP/PARAFAC (CP) decomposition tactics to update pairs of atoms or even whole dictionaries at once. We note that, with the exception of [16], [20], [21] and [22], tensor methods choose not to exploit the separable structure directly. The STARK algorithm [23] implicitly enforces Kronecker structure on the dictionary by solving a regularized convex relaxation of the hard rank-1 tensor recovery problem while TeFDiL [22] is a factorization based approach which imposes explicitly the Kronecker structure and finds the small dictionaries  $D_i$ .

MOD [24] solves the DL problem by viewing Y = DXas a least squares problem (LS) where the variable is X in the representation stage, and D in the dictionary update stage. TMOD is the *n*-dimensional adaptation of the MOD [18], [25], [26] algorithm. Given tensor  $T \in \mathbb{R}^{I_1 \times \cdots \times I_n}$  and matrix  $M \in$  $\mathbb{R}^{J imes I_k}$ , let  $T_{(k)}$  be the mode-k matricization and let  $\otimes_k$  be the mode-k product such that  $(T \otimes_k M)_{(k)} = MT_{(k)}$ . TMOD writes the *n*-dimensional DL problem as  $Y = X \otimes_1 D_1 \otimes_2$  $D_2 \otimes_3 \cdots \otimes_n D_n$ , where Y is the tensor containing N samples of  $m_1 \times \cdots \times m_n$  dimensions each and  $D_i \in \mathbb{R}^{m_i \times n_i}$  is the dictionary associated to dimension *i*. The right-hand variables are updated one at a time by solving a large LS problem. Similarly, the K-SVD algorithm has also been extended to the tensor setting [26]. As convolutional dictionaries are able to capture local structure in image data, separable convolutional dictionary learning was also introduced recently [27], [28].

In the sparse representation stage, these algorithms usually employ the 2D-OMP algorithm [11]. The complexity reduction brought by the separable version and its equivalence to 1D OMP [3] is thoroughly demonstrated in [12] and [11].

In coordinate descent fashion, TKSVD [16] simultaneously updates atom pairs i, j while keeping the other atoms fixed which can lead to an increase in the overall error (1) as shown in [12]. PairAK-SVD [12] is a direct non-tensor AK-SVD adaptation to the separable scenario which modifies the residual to update only one atom at a time.

STARK [23], SuKro [9], and [22] also extended TMOD by writing D as a sum of Kronecker products. Thus, in (1) we would write  $D = \sum_{r=1}^{R} D_2^{(r)} \otimes D_1^{(r)}$ . Dictionary update is also residual based but it involves a rank-R problem solved via ADMM [29]. Earlier results in [20] are similar to SuKro and also propose a non-separable version based on CP decomposition. The latest in this line of work is a DL algorithm that learns sums of *R* Kronecker products of K = 2 terms at a time [30]. Finally, the work in [22] described an online approach to the separable dictionaries learning problem.

While we focused on algorithmic developments, we also mention theoretical efforts made to understand the behavior and limits of these procedures like the local identifiability of the Kronecker-structured dictionaries [22], [31], [32] and sample complexity analysis of these dictionaries [33]–[35] that was shown to be significantly lower than for unstructured dictionaries [36]. Recently, matrix factorization were used to provide guarantees for global optimality [37].

#### III. THE PROPOSED ALGORITHM

In this section, we describe the proposed solution to the separable dictionary learning problem. We aim to solve the problem in (1) in three steps: 1) with  $D_1$  and  $D_2$  fixed we update all  $X_k$  using the 2D-OMP algorithm (alternatively, other sparse recovery algorithms can be used in this step but we choose an OMP approach due to its advantageous numerical properties); 2) for  $D_1$  and all  $X_k$  fixed we compute the optimal  $D_2$ , the minimizer of the objective function in (1); and 3) analogous to the previous step, we find the optimal  $D_1$ . We choose 2D-OMP over 1D-OMP in order to keep memory usage low and avoid vectorization operations that need to take place during the 1D sparse recovery algorithms. In this paper, we do not focus on developing new 2D sparse approximation methods but our contribution lies in developing new, efficient, ways of finding the two dictionaries  $D_1$  and  $D_2$ . Our dictionary update rules are based on the following result and corollary.

Proposition 1 (Optimal  $D_2$  update): Let  $D_1$  be fixed and denote  $Z_k = D_1 X_k$  with all  $X_k$ s fixed, then the minimizer of the objective function in (1) is

$$(\boldsymbol{D}_{2}^{\star})^{T} = \boldsymbol{W}_{2}(\sum_{k=1}^{N} \boldsymbol{Z}_{k}^{T} \boldsymbol{Z}_{k})^{-1}(\sum_{k=1}^{N} \boldsymbol{Z}_{k}^{T} \boldsymbol{Y}_{k}).$$
(2)

*Proof.* Ignore for now the norm constraints and the objective function of (1) develops into

$$\sum_{k=1}^{N} \left\| \boldsymbol{Z}_{k} \boldsymbol{D}_{2}^{T} - \boldsymbol{Y}_{k} \right\|_{F}^{2} = \left\| (\sum_{k=1}^{N} \boldsymbol{Z}_{k}^{T} \boldsymbol{Z}_{k})^{1/2} \boldsymbol{D}_{2}^{T} - (\sum_{k=1}^{N} \boldsymbol{Z}_{k}^{T} \boldsymbol{Z}_{k})^{-1/2} (\sum_{k=1}^{N} \boldsymbol{Z}_{k}^{T} \boldsymbol{Y}_{k}) \right\|_{F}^{2} + C.$$
(3)

Here  $C = \operatorname{tr}(\sum_{k=1}^{N} Y_{k}^{T} Y_{k} - (\sum_{k=1}^{N} Y_{k}^{T} Z_{k})(\sum_{k=1}^{N} Z_{k}^{T} Z_{k})^{-1}$  $(\sum_{k=1}^{N} Z_{k}^{T} Y_{k}))$  does not depend on  $D_{2}$ . Since the minimization above reduces to a standard least squares problem, the minimizer (assuming  $\sum_{k=1}^{N} Z_{k}^{T} Z_{k}$  has full rank, which happens almost always as  $N \gg m$ ) is  $(D_{2}^{\star})^{T} = (\sum_{k=1}^{N} Z_{k}^{T} Z_{k})^{-1} (\sum_{k=1}^{N} Z_{k}^{T} Y_{k})$ . Let  $W_{2}$  be a  $n_{2} \times n_{2}$  diagonal matrix such that  $D_{2}^{\star} W_{2}$  has normalized columns, then the optimal updates are  $D_{2}^{\star} \leftarrow D_{2}^{\star} W_{2}$  and  $X_{k} \leftarrow X_{k} W_{2}^{-1}$ ,

Algorithm 1 Distributed Separable Dictionary Learning

**Require:**  $Y \in \mathbb{R}^{m \times m \times N}$ ,  $D_{1,2} \in \mathbb{R}^{m \times n_{1,2}}$ , and  $s, p, K \in \mathbb{N}^*$  **Result:**  $D_1, D_2$ , and  $X \in \mathbb{R}^{n_1 \times n_2 \times N}$  (distributed at the nodes) **Setup:** Split as evenly as possible the *N* data points  $\{Y_i\}_{i=1}^p$ and distribute initial  $D_{1,2}$  among the *p* processing nodes **Main loop, for** 1,..., K:

**1.** Each node *i* computes the *s*-sparse representations  $X_k$  for all its data points  $Y_k$  and its partial sums  $P_i = \sum_k T_k T_k^T$  and  $R_i = \sum_k T_k Y_k^T$ ;

2. Master node accumulates  $P_i$  and  $R_i$  from all nodes *i*, sums  $P = \sum_{i=1}^{p} P_i$  and  $R = \sum_{i=1}^{p} R_i$ , and then computes  $D_1 = P^{-1}RW_1$ , according to Corollary 1. Master node distributes  $D_1$  to all *p* nodes;

3. With the freshly received  $D_1$ , each node *i* computes the *s*-sparse representations  $X_k$  for all its data points  $Y_k$  and its partial sums  $M_i = \sum_k Z_k^T Z_k$  and  $N_i = \sum_k Z_k^T Y_k$ ;

4. Master node accumulates  $M_i$  and  $N_i$  from all nodes i, sums  $M = \sum_{i=1}^{p} M_i$  and  $N = \sum_{i=1}^{p} N_i$ , and then computes  $D_2 = M^{-1}NW_2$ , according to Result 1. Master node distributes  $D_2$  to all p nodes.

without affecting the sparsity pattern (the objective function is minimized with the LS solution and the normalization diagonal cancels in the products  $X_k(D_2^{\star})^T$ ).

*Corollary 1 (Optimal*  $D_1$  *update):* Analogously to Result 1, denoting and fixing  $T_k = X_k D_2^T$  we have that the minimizer of the objective function in (1) is given by  $D_1^{\star} = (\sum_{k=1}^N Y_k T_k^T) (\sum_{k=1}^N T_k T_k^T)^{-1} W_1$ , where  $W_1$  is a diagonal matrix of size  $n_1 \times n_1$  chosen such that  $D_1^{\star}$  has normalized columns and update  $X_k \leftarrow W_1^{-1} X_k$  for all k.

Denoting  $G_1 = D_1^T D_1$  and  $G_2 = D_2 D_2^T$ , we note that we will compute  $Z_k^T Z_k^T = X_k^T G_1 X_k$  and  $T_k^T T_k^T = X_k G_2 X_k^T$ . We describe the full proposed distributed alternating optimization method in Algorithm 1. All dictionary updates guarantee a monotonic decrease in the objective function value but the same is not true about 2D OMP, in general. Algorithm 1 is an explicit, efficient parallel/distributed implementation of the classic TMOD approach for separable dictionary learning. In Figure 1, we provide a thread communication diagram that shows the initialization process and one iteration of the proposed method. We describe a scenario with p nodes but, for simplicity, show the communication/computation behavior only for two nodes highlighting also the dimensions of the matrices that are transferred - which do not depend on the size of the dataset N. In general, dictionary learning algorithms can be trivially parallelized in the sparse approximation step, i.e., the N sparse solutions  $X_k$  are computed separately. The issue is that each  $\mathbf{X}_k$  of size  $m \times m$  needs to be transferred/copied to a central node for the dictionary update step. For our method, just four  $m \times m$  matrices for each processing node are transferred and therefore the communication cost is lowered considerably, i.e.,  $O(m^2p)$  instead of  $O(m^2N)$ . As we will show experimentally, this allows the overall proposed algorithm to scale very well with the number of processing nodes p.

Remark 1 (Learning orthonormal dictionaries): If we impose in (1) the additional constraint that we are learning orthonormal dictionaries, i.e.,  $n_1 = n_2 = m$ ,  $D_1^T D_1 = D_1 D_1^T = I_m$  and  $D_2^T D_2 = D_2 D_2^T = I_m$ , then the least squares problem in Result 1 becomes an orthogonal Procrustes problem [38] whose solution is given by a singular value decomposition (for both  $D_1$  and  $D_2$ ). We call Algorithm 1 - Orthonormal, the same approach as Algorithm 1 but with the Procrustes updates for the dictionaries. Orthonormal dictionaries also improve the numerical complexity of the algorithm. The update formulas for the dictionaries become:  $D_1^* = U_1 V_1^T$  where  $(\sum_{k=1}^N Y_k^T D_2 X_k)(\sum_{k=1}^N X_k X_k^T) = U_1 \Sigma_1 V_1^T$  is given by the singular value decomposition (SVD) and  $D_2^* = U_2 V_2^T$  where  $(\sum_{k=1}^N Y_k^T D_1 X_k)(\sum_{k=1}^N X_k^T X_k) = U_2 \Sigma_2 V_2^T$  again by the SVD. Furthermore, 2D-OMP is no longer necessary as the sparse representations are computed as  $X_k = \mathcal{T}_s(D_1^T Y_k D_2)$ , where  $\mathcal{T}_s$  is an operator that keeps only the *s* largest (in absolute value) entries of the input matrix.

Remark 2 (On the computational complexity of Algorithm 1): We will highlight in the results section that the proposed method has a running time competitive against previously proposed algorithms from the literature. For simplicity let us assume that  $n_1 = n_2 = m$ . We focus on the dictionary updates for a single iteration of Algorithm 1 (in total we perform K iterations). First, notice that computing  $D_1^{\star}$  and  $D_2^{\star}$  takes about  $\frac{26}{3}m^3 + 12smN + 8m^2N + 4m^3N$  operations: the first term includes the two Cholesky decompositions  $(\frac{2}{3}m^3)$ , then the four back-substitutions  $(4m^3)$  used to solve the least squares problems for symmetric positive definite matrices with multiple right-hand sides and the two matrix-matrix multiplication to compute  $G_1$  and  $G_2$  (4 $m^3$ ), the second term is the cost of building all  $Z_k^T Z_k$  and  $T_k T_k^T$  (8smN) and the cost of computing all the sparse and non-sparse matrix products  $(4smN \text{ and } 4m^3N, \text{ respectively})$  and then summing them up  $(8m^2N)$  needed in Result 1 and Corollary 1. Because in general that  $N \gg m$  we see that the computational complexity is dominated by the construction of the matrix product sums.

From a computational perspective, tensor methods (TMOD, TKSVD, etc.) generally use a form of CP decomposition with alternating LS. Note that computing the inverse in the vectorized case costs  $O(m^6)$  operations, while in the 2D separable case it only takes  $O(m^3)$ .

Furthermore, note that the proposed approach is highly parallelizable: computing the *N* matrix products summations is distributed among multiple computing units (either CPUs or GPUs) with minimal cross-communication (only a partial summation matrix of size  $m \times m$  needs to be communicated). The same holds for the calculations of representations  $X_k$  which can be done locally at each processing unit without the need of communicating them explicitly.

*Remark 3 (Generalization to n dimensions):* Our result can be easily extended to more than two dimensions. We earlier described how TMOD generalizes MOD in *n* dimensions. Our method also holds in *n* dimensions using the property of the mode-*k* product:  $Y = X \otimes_1 D_1 \otimes_2 D_2 \otimes_3 \cdots \otimes_n D_n \iff$ 



Fig. 1. Assuming the overall dataset is split in p parts, we show the thread communication diagram with a master and two of the processing nodes highlighting the distributive nature of the proposed algorithm. We explicitly show all the dimensions of the matrices being communicated.

 TABLE I

 Average running times (seconds) for methods in Figure 2.

Method	<i>n</i> =	= 8	<i>n</i> = 16		
Wethod	s = 6	s = 8	s = 6	s = 8	
proposed, general $(p = 1)$	21	22	33	38	
proposed, ortho $(p = 1)$	18	20	18	22	
PairAK-SVD	140	171	205	266	
SuKro	23	27	47	56	
TeFDiL	8	10	19	24	
AK-SVD	10	14	16	21	

 $Y_{(k)} = D_k X_{(k)} (D_n \otimes D_{n-1} \otimes \cdots \otimes D_1)^T$ . Denoting with *T* the fixed dictionaries in the parenthesis we arrive at Corollary 1.

# IV. Simulation results $^1$

In this section we showcase the results achieved by the proposed algorithms (orthonormal and general Kronecker dictionaries) and compare with the state of the art.

First, in Figures 2 and 3 we show the RMSE achieved for various patches sizes *m*, dictionary sizes *n*, and dataset sizes *N*. Table I shows the running times of these methods (proposed methods run on p = 1 nodes) averaged over 5 rounds. As we are learning a relatively low number of dictionary parameters (roughly O(mn)) we choose *N* on the order of  $10^3$  and  $10^4$  image patches. The proposed general approach equals the performance of PairAK-SVD, but at a fraction of

the computational cost (on average  $\times 7$  speedup as compared to PairAK-SVD and ×1.7 against SuKro). TeFDiL is always faster than the general approach but the results are on average slightly worse. We would like to mention that the dictionary update step of TeFDiL involves large matrix-matrix multiplication (both matrices of size  $m^2 \times N$ ) and solving linear systems of size  $m^2 \times m^2$  with N right-hand sides, both of which are highly parallelizable in Matlab (in fact, running Matlab with the flag -singleCompThread leads to a 33-50% increase in the running time of TeFDiL). Therefore, one might say that TeFDiL is already "half-paralellized" (sparse representations are still computed sequentially). While TeFDiL is easy to parallelize (in a multi-thread fashion), the dictionary update step needs all the sparse representations X explicitly in the same processing node, so the communication cost in a distributed system would be high, as compared with our proposed approach where only small matrices need to be communicated (see Remark 2). The proposed orthogonal approach performs worse (of course due to the orthogonality constraints) but it is even faster - note that we report running time for 100 iterations but the proposed orthogonal approach converges much faster, mostly before 20 iterations. For perspective, we show the error achieved by an unstructured dictionary via AK-SVD [39] which is faster than the proposed method in its serial (p = 1) implementation. We now explore the speed-up benefits of the proposed method.

We have touted the efficient parallel implementation of Algorithm 1 and now in Table II we show the speedups achieved for both proposed approaches for various  $m \times m$  patches. The datasets, for each  $m \in \{8, 16, 32\}$ , have N = 128000,

<sup>&</sup>lt;sup>1</sup>Matlab & Python code at https://github.com/pirofti/ParallelSeparableDL



Fig. 2. Mean and standard deviation RMSE over 100 iterations with 5 random realizations of datasets. From left to right: 1) each dataset consists of N = 9216 randomly chosen non-overlapped  $8 \times 8$  image patches from a set of 12288 patches, n = 8 and s = 6; 2) analogous to 1) for s = 8; 3) analogous to 1) for n = 16 and s = 6; 4) analogous to 3) for s = 8.



Fig. 3. Mean and standard deviation RMSE against the size *n* of the dictionaries over 5 random realizations of datasets. From left to right: 1) each dataset consists of N = 9216 randomly chosen non-overlapped  $8 \times 8$  image patches from a set of 12288 patches, s = 6; 2) analogous to 1) for s = 8; 3) each dataset consists of N = 2304 randomly chosen non-overlapped  $16 \times 16$  image patches from a set of 3072 patches, s = 8; 4) analogous to 3) for s = 16.

 TABLE II

 Speedup achieved by the parallel implementation of the proposed algorithms (m = n). For ortho s = m, general s = 8.

Mathad		Number of processors $(p)$								
Method	m	1	2	4	6	8	12			
1	8	217 sec	$\times 2$	×4.3	×5.6	$\times 8$	×10.3			
proposed	16	86 sec	$\times 2$	×3.9	×5.4	×7.8	$\times 10.8$			
ortho	32	46 sec	$\times 2$	×3.8	×5.1	×7.1	×10.6			
proposed	8	440 sec	×0.6	×1.1	×1.4	$\times 2$	×2.6			
gaparal	16	450 sec	×1.5	×3.1	×3.8	×5.1	×5.4			
general	32	469 sec	$\times 2$	×3.4	$\times 4$	×4.6	$\times 5$			

N = 32000, and N = 8000 patches, respectively. For the proposed orthogonal case we scale almost precisely with the number of cores while for the proposed general case the speedup is less impressive. The latter observation is also due to the OMP implementation from the Python scikit-learn toolkit, which covers almost all the running time of the algorithm. We tested on an Intel(R) Xeon(R) CPU E5-2630 v4 @ 2.20GHz with 16 GB of RAM and 16 cores.

In Table III we reproduced the denoising experiments described in [12] with the proposed methods: we use  $512 \times 512$  images from the USC-SIPI database [40] and denoise via error

driven OMP with  $\varepsilon = 1.15\sigma_{\text{noise}}\sqrt{m}$ ; if the target error is not met, OMP is stopped when sparsity s = m/2 is reached. The results are averaged over 5 realizations of noise. We recorded negligible differences between realizations. This is standard practice in the literature [41].

For training we used N = 4000 patches of  $8 \times 8$  pixels with which we learned m = 8 by n = 16 dictionaries with target sparsity s = 6. The chosen images and noise levels are often used in the literature [8], [40], [41]. Denoising is performed on all N = 255025 overlapping patches and results are compared to the original image and measured in terms of peak signal-tonoise ratio (PSNR) and structural similarity index (SSIM) [42].

Finally, in Table IV we show DL execution times and speedups for the proposed methods on the Indian Pine 220-layered  $614 \times 1848$  hyperspectral image from [43]. The layers are divided among the *p* CPUs (in total *N* = 3913140 image patches) and the *m* = 8 by *n* = 8 dictionaries are trained for 100 iterations with *s* = 16. Even in this large scale experiment we observe scaling behavior similar to the results in Table II.

## V. CONCLUSIONS

In this paper we proposed two highly parallelizable and distributed algorithms for separable dictionary learning based

 TABLE III

 Average denoising PSNR(dB) and SSIM for standard images over 5 realizations. Best results are in boldface (PSNR) or underlined (SSIM).

	Method	lena		barbara		boat		peppers		house	
O noise / FSINK		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
5 / 34.1557	m = 8, n = 16, general	37.724	0.9358	33.045	0.9376	34.017	0.9031	36.372	0.9059	37.384	0.9277
	m = n = 8, ortho	38.381	0.9407	37.895	0.9611	36.977	0.9352	37.354	0.9194	38.904	0.9459
	PairAK-SVD [12]	37.912	0.9367	37.118	0.9588	35.987	0.9210	36.420	0.9066	38.038	0.9380
10 / 28.1322	m = 8, n = 16, general	35.132	0.9041	31.126	0.9044	32.276	0.8622	34.367	0.8705	34.912	0.8898
	m = n = 8, ortho	35.237	0.9048	33.949	0.9280	33.334	0.8738	34.553	0.8730	35.316	0.8944
	PairAK-SVD [12]	35.185	0.9046	33.720	0.9260	33.183	0.8724	34.4076	0.8708	35.116	0.8927
20 / 22.1105	m = 8, n = 16, general	31.988	0.8557	28.387	0.8400	29.723	0.7858	31.915	0.8323	32.275	0.8532
	m = n = 8, ortho	31.911	0.8548	29.833	0.8599	29.832	0.7857	31.836	0.8313	32.096	0.8507
	PairAK-SVD [12]	32.006	<u>0.8563</u>	30.027	<u>0.8634</u>	29.980	<u>0.7906</u>	31.930	<u>0.8323</u>	32.308	<u>0.8535</u>
30 / 18.5868	m = 8, n = 16, general	29.877	0.8174	26.921	0.7823	27.886	0.7252	30.116	0.8031	30.159	0.8211
	m = n = 8, ortho	29.848	0.8165	27.396	0.7909	27.834	0.7230	30.055	0.8021	30.012	0.8186
	PairAK-SVD [12]	29.937	0.8182	27.637	0.7973	27.989	0.7281	30.093	0.8024	30.250	0.8220
50 / 14.1505	m = 8, n = 16, general	27.358	0.7552	24.211	0.6695	25.486	0.6410	27.570	0.7531	27.296	0.7609
	m = n = 8, ortho	27.358	0.7550	24.449	0.6780	25.459	0.6397	27.513	0.7520	27.230	0.7592
	PairAK-SVD [12]	27.375	0.7557	24.604	0.6842	25.518	0.6421	27.559	0.7527	27.302	0.7614

TABLE IV Speedups achieved for Hyperspectral Experiments.

Number of processors ( <i>p</i> )									
Method	1	4	8	10	12	16			
ortho	20532 sec	×3.8	×7.6	×9.4	×10.8	×12.4			
general	325030 sec	×2.1	×5.5	$\times 8$	×9.1	×10.5			

on least squares dictionary updates. We show experimentally that the algorithms scale excellently with the number of cores or processing nodes and are competitive with the current state of the art separable dictionary learning methods for sparse representations. This approach opens the possibility of learning dictionaries given hundreds of thousands or millions of training signals in a reasonable amount of time.

#### References

- B. Dumitrescu and P. Irofti, Dictionary Learning Algorithms and Applications, Springer, 2018.
- [2] M. Elad, Sparse and Redundant Representations: from Theory to Applications in Signal Processing, Springer, 2010.
- [3] Y. C. Pati, R. Rezaiifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in 27th Asilomar Conf. Signals Systems Computers, Nov. 1993, vol. 1, pp. 40–44.
- [4] R. Rubinstein, M. Zibulevsky, and M. Elad, "Efficient implementation of the K-SVD algorithm using batch orthogonal matching pursuit," Tech. Rep. CS-2008-08, Technion Univ., Haifa, Israel, 2008.
- [5] A.-K. Seghouane and A. Iqbal, "Consistent adaptive sequential dictionary learning," Signal Processing, vol. 153, pp. 300–310, 2018.
- [6] Y. Yankelevsky and M. Elad, "Dual graph regularized dictionary learning," *IEEE Transactions on Signal and Information Processing* over Networks, vol. 2, no. 4, pp. 611–624, 2016.
- [7] A. Băltoiu, A. Pătraşcu, and P. Irofti, "Graph anomaly detection using dictionary learning," in *The 21st World Congress of the International Federation of Automatic Control*, 2020, pp. 1–8.
- [8] S. Hawe, M. Seibert, and M. Kleinsteuber, "Separable dictionary learning," in *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition, 2013, pp. 438–445.

- [9] C. F. Dantas, M. N. da Costa, and R. da Rocha Lopes, "Learning dictionaries as a sum of Kronecker products," *IEEE Signal Processing Letters*, vol. 24, no. 5, pp. 559–563, 2017.
- [10] S. Soltani, M. E. Kilmer, and P. C. Hansen, "A tensor-based dictionary learning approach to tomographic image reconstruction," *Bit Numer. Math.*, vol. 56, no. 4, pp. 1425–1454, 2016.
- [11] Y. Fang, J. Wu, and B. Huang, "2D sparse signal recovery via 2D orthogonal matching pursuit," *Science China Information Sciences*, vol. 55, no. 4, pp. 889–897, 2012.
- [12] P. Irofti and B. Dumitrescu, "Pairwise Approximate K-SVD," in 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2019, pp. 3677–3681.
- [13] P. Irofti and B. Dumitrescu, "GPU parallel implementation of the Approximate K-SVD algorithm using OpenCL," in 22nd European Signal Processing Conference, Lisbon, Portugal, 2014, pp. 271–275.
- [14] H. Raja and W. U. Bajwa, "Cloud K-SVD: A collaborative dictionary learning algorithm for big, distributed data," *IEEE Transactions on Signal Processing*, vol. 64, no. 1, pp. 173–188, 2016.
- [15] W. Dai, T. Xu, and W. Wang, "Simultaneous codeword optimization (SimCO) for dictionary update and learning," *IEEE Transactions on Signal Processing*, vol. 60, no. 12, pp. 6340–6353, 2012.
- [16] Y. Fu, J. Gao, Y. Sun, and X. Hong, "Joint multiple dictionary learning for tensor sparse coding," in *Neural Networks (IJCNN), 2014 International Joint Conference on*, 2014, pp. 2957–2964.
- [17] N. Qi, Y. Shi, X. Sun, J. Wang, B. Yin, and J. Gao, "Multi-dimensional sparse models," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 1, pp. 163–178, 2018.
- [18] K. Skretting, "Sparse signal representation using overlapping frames," Dr. ingeniøravhandling, 0809-103X; 2002:91, PhD. Thesis, 2002.
- [19] Y. Rivenson and A. Stern, "Compressed imaging with a separable sensing operator," *IEEE Signal Processing Letters*, vol. 16, no. 6, pp. 449–452, 2009.
- [20] S. H. Hsieh, C. S. Lu, and S. C. Pei, "2D sparse dictionary learning via tensor decomposition," in *Signal and Information Processing* (*GlobalSIP*), 2014 IEEE Global Conference on, 2014, pp. 492–496.
- [21] Z. Zhang and S. Aeron, "Denoising and completion of 3D data via multidimensional dictionary learning," in *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI-*16), 2016, pp. 2371–2377.
- [22] M. Ghassemi, Z. Shakeri, A. D. Sarwate, and W. U. Bajwa, "Learning mixtures of separable dictionaries for tensor data: Analysis and algorithms," *IEEE Transactions on Signal Processing*, vol. 68, pp. 33–48, 2020.
- [23] M. Ghassemi, Z. Shakeri, A. D. Sarwate, and W. U. Bajwa, "STARK: Structured dictionary learning through rank-one tensor recovery," in

2017 IEEE 7th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP), 2017, pp. 1–5.

- [24] K. Engan, S. O. Aase, and J. H. Husoy, "Method of optimal directions for frame design," in *IEEE Int. Conf. Acoustics Speech Signal Proc.*, 1999, vol. 5, pp. 2443–2446.
- [25] C. F. Caiafa and A. Cichocki, "Multidimensional compressed sensing and their applications," WIREs Data Mining and Knowledge Discovery, vol. 3, no. 6, pp. 355–380, 2013.
- [26] F. Roemer, G. Del Galdo, and M. Haardt, "Tensor-based algorithms for learning multidimensional separable dictionaries," in 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2014, pp. 3963–3967.
- [27] J. Quesada, P. Rodriguez, and B. Wohlberg, "Separable dictionary learning for convolutional sparse coding via split updates," in 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2018, pp. 4094–4098.
- [28] G. Silva, J. Quesada, and P. Rodriguez, "Efficient separable filter estimation using rank-1 convolutional dictionary learning," in 2018 IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP), 2018, pp. 1–6.
- [29] S. Boyd, N. Parikh, E. Chu, B. Peleato, J. Eckstein, et al., "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [30] C. F. Dantas, J. E. Cohen, and R. Gribonval, "Learning fast dictionaries for sparse representations using low-rank tensor decompositions," in *Latent Variable Analysis and Signal Separation*, Y. Deville, S. Gannot, R. Mason, M. D. Plumbley, and D. Ward, Eds. 2018, pp. 456–466, Springer International Publishing.
- [31] Z. Shakeri, A. D. Sarwate, and W. U. Bajwa, "Identifiability of Kronecker-structured dictionaries for tensor data," *IEEE Journal of Selected Topics in Signal Processing*, vol. 12, no. 5, pp. 1047–1062, 2018.
- [32] M. Ghassemi, Z. Shakeri, W. U. Bajwa, and A. D. Sarwate, "Sample complexity bounds for low-separation-rank dictionary learning," in 2019 IEEE International Symposium on Information Theory (ISIT), 2019, pp. 2294–2298.
- [33] Z. Shakeri, W. U. Bajwa, and A. D. Sarwate, "Minimax lower bounds for Kronecker-structured dictionary learning," in 2016 IEEE International Symposium on Information Theory (ISIT), 2016, pp. 1148–1152.
- [34] Z. Shakeri, W. U. Bajwa, and A. D. Sarwate, "Sample complexity bounds for dictionary learning of tensor data," in 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2017, pp. 4501–4505.
- [35] Z. Shakeri, A. D. Sarwate, and W. U. Bajwa, "Sample complexity bounds for dictionary learning from vector- and tensor-valued data," in *Information-Theoretic Methods in Data Science*, M. Rodrigues and Y. C. Eldar, Eds., pp. 134–162. Cambridge University Press, Cambridge, UK, 2019.
- [36] A. Jung, Y. C. Eldar, and N. Görtz, "On the minimax risk of dictionary learning," *IEEE Transactions on Information Theory*, vol. 62, no. 3, pp. 1501–1515, 2016.
- [37] E. Schwab, B. D. Haeffele, R. Vidal, and N. Charon, "Global optimality in separable dictionary learning with applications to the analysis of diffusion MRI," *SIAM Journal on Imaging Sciences*, vol. 12, no. 4, pp. 1967–2008, 2019.
- [38] P. Schonemann, "A generalized solution of the orthogonal Procrustes problem," *Psychometrika*, vol. 31, no. 1, pp. 1–10, 1966.
- [39] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Proc.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [40] A. G. Weber, "The USC-SIPI Image Database," 1997.
- [41] M. Elad and M. Aharon, "Image denoising via sparse and redundant representations over learned dictionaries," *IEEE Transactions on Image Processing*, vol. 15, no. 12, pp. 3736–3745, 2006.
- [42] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Proc.*, vol. 13, no. 4, pp. 600–612, 2004.
- [43] M. F. Baumgardner, L. L. Biehl, and D. A. Landgrebe, "220 Band AVIRIS Hyperspectral Image Data Set: June 12, 1992 Indian Pine Test Site 3," Sep 2015.