

Robust Kronecker Component Analysis

Mehdi Bahri, *Student Member, IEEE*, Yannis Panagakis, and Stefanos Zafeiriou, *Member, IEEE*

Abstract—Dictionary learning and component analysis models are fundamental for learning compact representations that are relevant to a given task (feature extraction, dimensionality reduction, denoising, etc.). The model complexity is encoded by means of specific structure, such as sparsity, low-rankness, or nonnegativity. Unfortunately, approaches like K-SVD - that learn dictionaries for sparse coding via Singular Value Decomposition (SVD) - are hard to scale to high-volume and high-dimensional visual data, and fragile in the presence of outliers. Conversely, robust component analysis methods such as the Robust Principal Component Analysis (RPCA) are able to recover low-complexity (e.g., low-rank) representations from data corrupted with noise of unknown magnitude and support, but do not provide a dictionary that respects the structure of the data (e.g., images), and also involve expensive computations. In this paper, we propose a novel Kronecker-decomposable component analysis model, coined as Robust Kronecker Component Analysis (RKCA), that combines ideas from sparse dictionary learning and robust component analysis. RKCA has several appealing properties, including robustness to gross corruption; it can be used for low-rank modeling, and leverages separability to solve significantly smaller problems. We design an efficient learning algorithm by drawing links with a restricted form of tensor factorization, and analyze its optimality and low-rankness properties. The effectiveness of the proposed approach is demonstrated on real-world applications, namely background subtraction and image denoising and completion, by performing a thorough comparison with the current state of the art.

Index Terms—Component analysis, dictionary learning, separable dictionaries, low-rank, sparsity, global optimality

1 INTRODUCTION

COMPONENT analysis models and representation learning methods are flexible data-driven alternatives to analytical dictionaries (e.g., Fourier analysis, or wavelets) for signal and data representation. The underlying principle is to solve an optimization problem that encodes the desired properties of the representations and of the bases, and describes what task(s) the representation should help solving. Starting from Principal Component Analysis [1], [2], a rich set of algorithms has been developed for feature extraction, dimensionality reduction, clustering, classification, or denoising—to name but a few. The importance of learned components and representations cannot be overstated, and neither can their effectiveness in dramatically improving machine perception. Prominent examples are Convolutional Neural Networks [3], [4], which through hierarchical feature extraction build ad-hoc representations enabling state of the art performance on a wide range of problems [5].

In this work, we propose the Robust Kronecker Component Analysis (RKCA) family of algorithms for the unsupervised learning of compact representations of tensor data.

Our method offers to bridge (multilinear) Robust PCA [6], [7] and Sparse Dictionary Learning [8], [9] from the perspective of a robust low-rank tensor factorization. Although our method is generic enough to be presented for arbitrary tensors, we focus on 3-dimensional tensors, and especially those obtained by concatenating 2-dimensional tensor observations (i.e., matrices). We present a framework for jointly learning a (Kronecker) separable dictionary and sparse representations in the presence of outliers, such that the learned dictionary is low-rank, and the outliers are separated from the data. The double perspective adopted allows us to draw on recent work from the tensor factorization literature to provide some theoretical optimality guarantees, discussed in Section 4.

1.1 Robust PCA and Sparse Dictionary Learning

Assuming a set of N data samples $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^m$ represented as the columns of a matrix \mathbf{X} , structured matrix factorizations seek to decompose \mathbf{X} into meaningful components of a given structure, by solving a regularization problem of the form:

$$\min_{\mathbf{Z}} l(\mathbf{X}, \mathbf{Z}) + \lambda g(\mathbf{Z}), \quad (1)$$

where \mathbf{Z} is an approximation of the data with respect to a loss $l(\cdot)$, $g(\cdot)$ is a possibly non-smooth regularizer that encourages the desired structure, and $\lambda \geq 0$ is a regularization parameter balancing the two terms. Popular instances of Equation (1) include Principal Component Analysis (PCA) [1], [2] and its variants, e.g., Sparse PCA [10], Robust PCA (RPCA) [11], as well as sparse dictionary learning [8], [9], [12].

Concretely, when \mathbf{Z} is taken to be factorized as $\mathbf{Z} = \mathbf{D}\mathbf{R}$ we obtain a range of different models depending on the choice of the regularization and of the properties of \mathbf{D} . For instance,

- M. Bahri is with the Department of Computing, Imperial College London, London SW7 2RH, United Kingdom. E-mail: mehdi.bahri15@imperial.ac.uk.
- Y. Panagakis is with the Department of Computing, Imperial College London, London SW7 2RH, United Kingdom, and also with Middlesex University, London NW4 4BT, United Kingdom. E-mail: i.panagakis@imperial.ac.uk.
- S. Zafeiriou is with the Department of Computing, Imperial College London, London SW7 2RH, United Kingdom, and also with the University of Oulu, Oulu 90014, Finland. E-mail: s.zafeiriou@imperial.ac.uk.

Manuscript received 15 Jan. 2018; revised 1 Oct. 2018; accepted 6 Nov. 2018. (Corresponding author: Mehdi Bahri.)

Recommended for acceptance by L. Liu, M. Pietkainen, J. Chen, G. Zhao,

X. Wang, and R. Chellappa.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TPAMI.2018.2881476

assuming $\mathbf{Z} = \mathbf{D}\mathbf{R}$ is a low-rank approximation of \mathbf{X} and $\lambda = 0$, Equation (1) yields PCA, while by imposing sparsity on \mathbf{D} , Sparse PCA [10] is obtained. To handle data corrupted by sparse noise of large magnitude, RPCA [11] assumes that the observation matrix, \mathbf{X} , is the sum of a low-rank matrix \mathbf{A} and of a sparse matrix \mathbf{E} that collects the gross errors, or outliers. This model is actually a special instance of Equation (1) when $\mathbf{Z} = \mathbf{A} + \mathbf{E}$, $g(\mathbf{Z}) = \|\mathbf{A}\|_* + \lambda\|\mathbf{E}\|_1$, and $l(\cdot)$ is the Frobenius norm. Here, $\|\cdot\|_*$ denotes the low-rank promoting nuclear norm and $\|\cdot\|_1$ denotes the ℓ_1 norm that enforces sparsity. Matrix RPCA has been extended to tensors in multiple ways, relying on varying definitions of the tensor rank. We refer to [6], [7] for an overview of CP-based and Tucker-based tensor RPCA models, and to Section 7 for specifics on the models compared in this paper.

Assuming \mathbf{D} is over-complete and requiring $\mathbf{R} = [\mathbf{r}_1, \dots, \mathbf{r}_n]$ to be sparse, Equation (1) leads to sparse dictionary learning by solving the non-convex optimization problem:

$$\min_{\mathbf{R}, \mathbf{D}} \|\mathbf{X} - \mathbf{D}\mathbf{R}\|_F^2 + \lambda \sum_{i=1}^n \|\mathbf{r}_i\|_0, \quad (2)$$

where $\|\cdot\|_F$ is the Frobenius norm and $\|\cdot\|_0$ is the ℓ_0 pseudo-norm, counting the number of non-zero elements. In K-SVD and its variants, problem (2) is solved in an iterative manner that alternates between sparse coding of the data samples on the current dictionary, and a process of updating the dictionary atoms to better fit the data using the Singular Value Decomposition (SVD) [13], [14].

1.2 Limits of the Classical Approach

While the Robust PCA and the sparse dictionary learning paradigms have been immensely successful in practice, they suffer from limitations that can restrict their applicability.

The recovery guarantees of Robust PCA and of similar compressed-sensing approaches have been derived under strong incoherence assumptions, that may not be satisfied even if the data strictly follows the low-rank assumption. In fact, whenever the data lies on a low-dimensional linear subspace but isn't uniformly distributed, the *extra structure* can induce significant coherence. [15] studies this problem in details, and in [15], [16], the authors show the adverse effects of coherence can be mitigated in existing low-rank modeling methods, such as [17], by choosing the representation dictionary to be low-rank.

Modern representation learning methods have to deal with increasingly large amounts of increasingly high-dimensional data. Classical low-rank modeling and dictionary learning models tend to be too expensive for that setting: typically, algorithms of the K-SVD family suffer from a high computational burden, preventing their applicability to high-dimensional and large scale data.

To overcome the issue of scalability in dictionary learning, a separable structure on the dictionary can be enforced. For instance, Separable Dictionary Learning (SeDiL) [18] considers a set of samples in matrix form, namely, $\mathcal{X} = (\mathbf{X}_i)_i$, admitting sparse representations on a pair of bases \mathbf{A}, \mathbf{B} , of which the Kronecker product constructs the dictionary. The corresponding objective is:

$$\min_{\mathbf{A}, \mathbf{B}, \mathcal{R}} \frac{1}{2} \sum_i \|\mathbf{X}_i - \mathbf{A}\mathbf{R}_i\mathbf{B}^T\|_F^2 + \lambda g(\mathcal{R}) + \kappa r(\mathbf{A}) + \kappa r(\mathbf{B}), \quad (3)$$

where the regularizers $g(\cdot)$ and $r(\cdot)$ promote sparsity in the representations, and low mutual-coherence of the dictionary $\mathbf{D} = \mathbf{B} \otimes \mathbf{A}$, respectively. Here, \mathbf{D} is constrained to have orthogonal columns, i.e., the pair \mathbf{A}, \mathbf{B} shall lie on the product manifold of two product of sphere manifolds. A different approach is taken in [19]: a separable 2D dictionary is learned in a two-step strategy similar to that of K-SVD. Each matrix observation \mathbf{X}_i is represented as $\mathbf{A}\mathbf{R}_i\mathbf{B}^T$. In the first step, the sparse representations \mathbf{R}_i are found by 2D Orthogonal Matching Pursuit (OMP) [20]. In the second step, a CANDECOMP/PARAFAC (CP) [21], [22], [23] decomposition is performed on a tensor of residuals via Regularized Alternating Least Squares to solve $\min_{\mathbf{A}, \mathbf{B}, \mathcal{R}} \|\mathcal{X} - \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B}\|_F$.¹ However, those models learn over-complete dictionaries on a multitude of small patches extracted from the images. For large images or large number of images, the number of patches can easily become prohibitively large, undermining the scalability benefits of learning separable dictionaries. Moreover, none of these models is robust to gross corruption in the dataset.

Beyond scalability, separable dictionaries have been shown to be theoretically appealing when dealing with tensor-valued observations. According to [24], the necessary number of samples for accurate (up to a given error) reconstruction of a Kronecker-structured dictionary within a local neighborhood scales with the sum of the product of the dimensions of the constituting dictionaries when learning on tensor data (see [25] for 2-dimensional data, and [26], [27] for N -order tensor data), compared to the product for vectorized observations. This suggests better performance is achievable compared to classical methods on tensor observations.

1.3 Outline and Contributions

Here, we propose novel methods for separable dictionary learning based on robust tensor factorisations that learn simultaneously the dictionary and the sparse representations. We do not seek overcompleteness, but rather promote low-rankness in a pair of dictionaries, and sparsity in the codes to learn a low-rank representation of the input tensor. In this regard, our methods combine ideas from both Sparse Dictionary Learning and Robust PCA, as well as tensor factorizations. Our solvers are based on the Alternating Direction of Multipliers Method (ADMM) [28].

A preliminary version of this work has been presented in [29]. This paper offers the following novelties:

- We generalize the results of [29] and propose new regularizers that yield stronger optimality properties, we call the resulting framework Robust Kronecker Component Analysis.
- Unlike previous models that rely solely on regularization to impose low-rankness, our work presents surrogates of specific definitions of the tensor rank that impose low-rankness through both regularisation and structure. Specifically, our method can be seen as modeling a low-rank tensor by tensor sparse coding with additional regularization on the factors.
- We show that RKCA with these well-chosen regularizers can be reformulated in the framework of [30], [31] allowing us to provide global optimality

1. cf. Definition 2.4 for the product \times_n

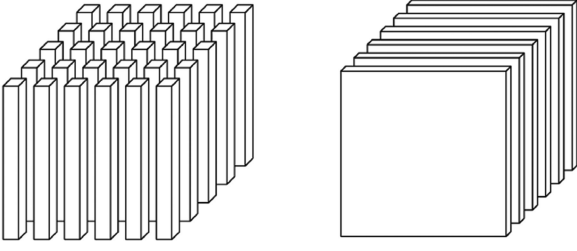


Fig. 1. Mode-1 fibers (tubes) and mode-3 slices of a 3-way tensor. Adapted from [33].

guarantees. It is worth mentioning that our proof applies to any Tucker factorization problem.

- We demonstrate that RKCA can perform tensor completion in the presence of gross corruption.
- We derive a Linearized ADMM (LADMM) algorithm for RKCA to improve scalability.
- The experimental evaluation has been enriched to include the Low-Rank Representation method [17], as well as comparisons with the recent Deep Image Prior [32].
- Finally, we offer new perspectives on the low-rank promoting properties of RKCA.

To the best of our knowledge, our method is the first to leverage dictionaries that are both low-rank and separable for robust representation learning.

The rest of the manuscript is organized as follows. Section 2 reviews fundamental definitions and results of tensor algebra and of the Kronecker product. Section 3 is dedicated to deriving the RKCA model, relating RKCA to separable dictionary learning, and deriving RKCA with missing values. In Section 4, we discuss optimality guarantees by formulating RKCA as an equivalent CP factorization with duplicated factors. Perspectives on the low-rank promoting properties can be found in Section 5, and a discussion on the computational cost and implementation details of the methods in Section 6. Finally, we present in Section 7 experimental evidence of the effectiveness of RKCA on synthetic and real-world data.

2 PRELIMINARIES

In this section we review fundamental properties of tensor algebra and of the Kronecker product, and present the notations and conventions followed in the paper.

2.1 Tensor Algebra

We refer to multidimensional arrays of real numbers of dimension $I_1 \times I_2 \times \dots \times I_N$ as N -dimensional, or N -way real tensors. The order of the tensor is the number of indices required to address its elements. Consequently, each element of an N th-order tensor \mathcal{X} is addressed by N indices, i.e., $(\mathcal{X})_{i_1, i_2, \dots, i_N} \doteq x_{i_1, i_2, \dots, i_N}$. We denote tensors by bold calligraphic letters, e.g., \mathcal{X} .

The sets of real and integer numbers are denoted by \mathbb{R} and \mathbb{Z} , respectively. An N th-order real-valued tensor \mathcal{X} is defined over the tensor space $\mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$, where $I_n \in \mathbb{Z}$ for $n = 1, 2, \dots, N$.

Matrices (vectors) are second order (first order) tensors and are denoted by uppercase (lowercase) bold letters, e.g., \mathbf{X} (\mathbf{x}). The i th column of a matrix \mathbf{X} will be written \mathbf{x}_i for convenience.

2.1.1 Elementary Definitions

Tensor fibers and tensor slices are special subsets of the elements of a tensor.

Definition 2.1 (Tensor slices and fibers). We define the mode- n tensor fibers as all subsets of the tensor elements obtained by fixing all but the n th index to given values. Similarly, tensor slices are obtained by fixing all but two indices.

Fig. 1 illustrates the concepts of fibers and slices for a third order tensor. For a third order tensor, the slices are effectively matrices obtained by “slicing” the tensor along a mode.

A tensor’s elements can be re-arranged to form a new tensor of different dimension. We define the tensor *vectorisation* and tensor *matricisations* (or *unfoldings*) as re-arrangements of particular interest.

Definition 2.2 (Tensor vectorisation). Let $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ be an N th-order tensor. The tensor vectorisation, or flattening, operator $\mathcal{X} \mapsto \text{vec}(\mathcal{X})$ maps each element x_{i_1, i_2, \dots, i_N} of \mathcal{X} to a unique element x_j of a real-valued vector $\text{vec}(\mathcal{X})$ of dimension $\prod_{m=1}^N I_m$ with the following bijection:

$$j = 1 + \sum_{k=1}^N (i_k - 1)J_k \quad \text{and} \quad J_k = \prod_{m=1}^{k-1} I_m. \quad (4)$$

Definition 2.3 (Mode- n unfolding and tensorisation).

Let $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ be an N th-order tensor. The mode- n tensor matricisation - or unfolding, with $n \in \{1, 2, \dots, N\}$, is the matrix $\mathbf{X}_{[n]}$ of dimensions $(I_n, \prod_{k \neq n} I_k)$ such that tensor entry x_{i_1, i_2, \dots, i_N} is mapped to a unique element $x_{i_n, j}$, with the following bijection:

$$j = 1 + \sum_{\substack{k=1 \\ k \neq n}}^N (i_k - 1)J_k \quad \text{and} \quad J_k = \prod_{\substack{m=1 \\ m \neq n}}^{k-1} I_m. \quad (5)$$

The inverse operator of the mode- n unfolding is the mode- n tensorisation and is denoted $\text{fold}_n(\mathbf{X}_{[n]}) = \mathcal{X}$.

In other words, the mode- n unfolding is the matrix whose columns are the mode- n fibers obtained by first varying I_1 , then I_2 , up until I_N with the exception of I_n .

Having defined both tensor fibers, slices, and unfoldings; we can now define a higher-order analogue for matrix-vector multiplication as the *tensor mode- n product*.

Definition 2.4 (Mode- n product). The mode- n product of $\mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ with $\mathbf{U} \in \mathbb{R}^{J \times I_n}$ is the tensor $\mathcal{X} \times_n \mathbf{U} \in \mathbb{R}^{I_1 \times \dots \times I_{n-1} \times J \times I_{n+1} \times \dots \times I_N}$ such that

$$(\mathcal{X} \times_n \mathbf{U})_{i_1 \dots i_{n-1} j i_{n+1} i_N} = \sum_{i_n=1}^{I_n} x_{i_1 i_2 \dots i_N} u_{j i_n}. \quad (6)$$

Effectively, each mode- n fiber is multiplied by \mathbf{U} .

The mode- n product can be expressed as a matrix product as per Proposition 2.1.

Proposition 2.1. With the notations of Definition 2.4:

$$(\mathcal{X} \times_{i=1}^N \mathbf{U}_i)_{[n]} = \mathbf{U}_n \mathbf{X}_{[n]} (\otimes_{i=N, i \neq n}^1 \mathbf{U}_i)^\top, \quad (7)$$

where \otimes is the Kronecker product of Definition 2.5.

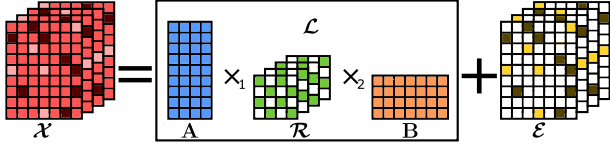


Fig. 2. Illustration of the decomposition.

2.1.2 Tensor Rank

Contrary to matrices, the tensor rank is not uniquely defined. In this paper, we are interested in the tensor *Tucker rank*, which we define as the vector of the ranks of a tensor's mode- n unfoldings (i.e., its mode- n ranks), and the tensor *multi-rank* [34], [35] defined as the vector of the ranks of its frontal slices.

More details about tensors, such as the definitions of the tensor CP-rank and information about common decompositions can be found in [33], for example.

2.2 Properties of the Kronecker Product

We now remind the reader of the definition and of important properties of the Kronecker product.

Definition 2.5 (Kronecker product). Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{B} \in \mathbb{R}^{p \times q}$ then

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{11}\mathbf{B} & a_{12}\mathbf{B} & \dots & a_{1n}\mathbf{B} \\ \vdots & \vdots & \vdots & \vdots \\ a_{n1}\mathbf{B} & a_{n2}\mathbf{B} & \dots & a_{nn}\mathbf{B} \end{bmatrix}. \quad (8)$$

Proposition 2.2 is a fundamental result relating the Kronecker product and tensor vectorisation.

Proposition 2.2. Let $\mathcal{Y} = \mathcal{X} \times_1 \mathbf{U}_1 \times_2 \mathbf{U}_2 \dots \times_N \mathbf{U}_N$ then

$$\text{vec}(\mathcal{Y}) = (\otimes_{i=1}^N \mathbf{U}_i) \text{vec}(\mathcal{X}). \quad (9)$$

In particular, if $\mathbf{Y} = \mathbf{A}\mathbf{X}\mathbf{B}^\top$ then

$$\text{vec}(\mathbf{Y}) = (\mathbf{B} \otimes \mathbf{A}) \text{vec}(\mathbf{X}). \quad (10)$$

The Kronecker product is *compatible* with most common matrix decomposition. For the Singular Value Decomposition, we have Proposition 2.3.

Proposition 2.3. Let $\mathbf{Y} = \mathbf{U}\mathbf{S}\mathbf{V}^\top$ and $\mathbf{X} = \mathbf{H}\mathbf{T}\mathbf{G}^\top$ then

$$\mathbf{Y} \otimes \mathbf{X} = (\mathbf{U} \otimes \mathbf{H})(\mathbf{S} \otimes \mathbf{T})(\mathbf{V} \otimes \mathbf{G})^\top. \quad (11)$$

Where we used the identity $(\mathbf{A} \otimes \mathbf{B})^\top = \mathbf{A}^\top \otimes \mathbf{B}^\top$.

We refer the reader to one of the many linear algebra references for further information about the Kronecker product and its properties.

3 MODEL DERIVATION

In this section we start by describing RKCA as a structured tensor factorization. We then show its equivalence to a dictionary learning problem, discuss how to extend the model to handle missing values, and finally derive a first algorithm that will serve as the basis for the rest of the discussion.

3.1 The Tensor Factorization Perspective

Consider a set of N two dimensional observations (i.e., matrices) $\mathbf{X}_i \in \mathbb{R}^{m \times n}$, $i = 1, \dots, N$ stacked as the frontal slices of a tensor \mathcal{X} . We study *Tensor Robust PCA* problems:

$$\begin{aligned} \min_{\mathcal{L}, \mathcal{E}} \quad & f(\mathcal{L}) + g(\mathcal{E}), \\ \text{s.t.} \quad & \mathcal{X} = \mathcal{L} + \mathcal{E}, \end{aligned} \quad (12)$$

where \mathcal{L} and \mathcal{E} are respectively the low-rank and sparse components of \mathcal{X} . We will define $f(\cdot)$ and $g(\cdot)$ to be regularization functions, possibly non-smooth and non-convex, meant to promote structural properties of \mathcal{L} and \mathcal{E} - in our case, low-rankness of \mathcal{L} , and sparsity in \mathcal{E} .

The Robust Kronecker Component Analysis is obtained by assuming \mathcal{L} factorizes in a restricted form of Tucker factorization, and defining $f(\cdot)$ as a combination of penalties on the factors. More specifically, we assume:

$$\mathcal{L} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B}. \quad (13)$$

Fig. 2 illustrates the decomposition.

We construct f to promote low-rankness in the first two modes of \mathcal{L} and therefore in each of its frontal slices. In this work, we will discuss three different choices of f depending on the *positive-homogeneity* degree of the full regularization.

Definition 3.1 (Def. 3 of [31, page 6]). A function $\theta : \mathbb{R}^{D^1} \times \dots \times \mathbb{R}^{D^N} \rightarrow \mathbb{R}^D$ is positively homogeneous with degree p if $\theta(\alpha x^1, \dots, \alpha x^N) = \alpha^p \theta(x^1, \dots, x^N)$, $\forall \alpha \geq 0$. Notably $\theta(0) = 0$ holds.

The following choice of f is the *degree 2 regularizer* and recovers the model presented in [29]:

$$f(\mathcal{L}) = \alpha \|\mathcal{R}\|_1 + \|\mathbf{B} \otimes \mathbf{A}\|_F, \quad (14)$$

while choice

$$f(\mathcal{L}) = \alpha \|\mathcal{R}\|_1 \|\mathbf{B} \otimes \mathbf{A}\|_F, \quad (15)$$

and

$$f(\mathcal{L}) = \alpha \|\mathcal{R}\|_1 \|\mathbf{B} \otimes \mathbf{A}\|_*, \quad (16)$$

are *degree 3 regularizers*.

In all three cases, the regularizer comprises a norm on the Kronecker product of the Tucker factors, hence the proposed method is coined as RKCA. The interpretation of this is made clear in Section 3.2.

3.2 Robust Separable Dictionary Learning

We now establish how to interpret RKCA as a robust separable dictionary learning scheme for tensor-valued observations. We remind the reader that we do not seek overcompleteness, but rather choose to learn low-rank dictionaries as a mean of modeling low-rankness in the observations. We impose separability to preserve the spatial correlation within the tensor inputs, and for scalability (c.f., Section 3.4). These choices are further motivated by the work of [15] (for low-rankness), and [24] (for separability).

Consider the following Sparse Dictionary Learning problem with Frobenius-norm regularization on the dictionary

\mathbf{D} , where we decompose N observations $\mathbf{x}_i \in \mathbb{R}^{mn}$ on $\mathbf{D} \in \mathbb{R}^{mn \times r_1 r_2}$ with representations $\mathbf{r}_i \in \mathbb{R}^{r_1 r_2}$:

$$\min_{\mathbf{D}, \mathbf{R}} \sum_i \|\mathbf{x}_i - \mathbf{D}\mathbf{r}_i\|_2^2 + \lambda \sum_i \|\mathbf{r}_i\|_1 + \|\mathbf{D}\|_F. \quad (17)$$

We assume a Kronecker-decomposable dictionary $\mathbf{D} = \mathbf{B} \otimes \mathbf{A}$ with $\mathbf{A} \in \mathbb{R}^{m \times r_1}$, $\mathbf{B} \in \mathbb{R}^{n \times r_2}$. To model the presence of outliers, we introduce a set of vectors $\mathbf{e}_i \in \mathbb{R}^{mn}$ and, with $d = r_1 r_2 + mn$, define the block vectors and matrices:

$$\mathbf{y}_i = \begin{bmatrix} \mathbf{r}_i \\ \mathbf{e}_i \end{bmatrix} \in \mathbb{R}^d \quad \mathbf{C} = [\mathbf{B} \otimes \mathbf{A} \quad \mathbf{I}] \in \mathbb{R}^{mn \times d}. \quad (18)$$

We obtain a two-level structured dictionary \mathbf{C} and the associated sparse encodings \mathbf{y}_i . Breaking-down the variables to reduce dimensionality and discarding $\|\mathbf{I}\|_F$:

$$\min_{\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathbf{E}} \sum_i \|\mathbf{x}_i - (\mathbf{B} \otimes \mathbf{A})\mathbf{r}_i - \mathbf{e}_i\|_2^2 + \lambda \sum_i \|\mathbf{r}_i\|_1 + \lambda \sum_i \|\mathbf{e}_i\|_1 + \|\mathbf{B} \otimes \mathbf{A}\|_F. \quad (19)$$

Suppose now that the observations \mathbf{x}_i were obtained by vectorizing two-dimensional data such as images, i.e., $\mathbf{x}_i = \text{vec}(\mathbf{X}_i)$, $\mathbf{X}_i \in \mathbb{R}^{m \times n}$. Without loss of generality, we choose $r_1 = r_2 = r$ and $\mathbf{r}_i = \text{vec}(\mathbf{R}_i)$, $\mathbf{R}_i \in \mathbb{R}^{r \times r}$, and recast the problem as:

$$\min_{\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathbf{E}} \sum_i \|\mathbf{X}_i - \mathbf{A}\mathbf{R}_i\mathbf{B}^\top - \mathbf{E}_i\|_F^2 + \lambda \sum_i \|\mathbf{R}_i\|_1 + \lambda \sum_i \|\mathbf{E}_i\|_1 + \|\mathbf{B} \otimes \mathbf{A}\|_F. \quad (20)$$

Equivalently, enforcing the equality constraints and concatenating the matrices \mathbf{X}_i , \mathbf{R}_i , and \mathbf{E}_i as the frontal slices of 3-way tensors, we obtain problem Equation (21).

$$\min_{\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathbf{E}} \alpha \|\mathbf{R}\|_1 + \lambda \|\mathbf{E}\|_1 + \|\mathbf{B} \otimes \mathbf{A}\|_F, \quad (21)$$

$$\text{s.t.} \quad \mathcal{X} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} + \mathcal{E},$$

where \mathcal{R} is a tensor whose i th frontal slice encodes the representation of \mathbf{X}_i on the column basis \mathbf{A} and on the row basis \mathbf{B} . Problem Equation (21) corresponds to choosing f to be Equation (14) and g to be the element-wise ℓ_1 norm in Equation (12).

We verify $r \in \mathbb{N}$, $r \leq \min(m, n)$ is a natural *upper bound* on the rank of each frontal slice of \mathcal{L} and on its mode-1 and mode-2 ranks:

$$\mathbf{L}_i = \mathbf{A}\mathbf{R}_i\mathbf{B}^\top. \quad (22)$$

We have thus shown how the two perspectives are equivalent.

3.3 Tensor Completion with RKCA

Handling outliers is important for real-world applications, but data corruption is polymorphic, and missing values are common. We now present an extension to the original problem where we assume we are only given incomplete and possibly grossly-corrupted observations, and wish to perform simultaneous removal of the outliers and imputation of the missing values.

Given a set $\Omega \subset \mathbb{N}^{I_1 \times I_2 \times \dots \times I_N}$ we define the sampling operator $\pi_\Omega : \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N} \rightarrow \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$ as the projection on the space of N -way tensors whose only non-zero entries are indexed by N -tuples in Ω , i.e., $\forall \mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$:

$$\pi_\Omega(x_{i_1, i_2, \dots, i_N}) = \begin{cases} 1 & (i_1, i_2, \dots, i_N) \in \Omega \\ 0 & (i_1, i_2, \dots, i_N) \notin \Omega \end{cases}.$$

It is simple to show π_Ω is an orthogonal projection. Clearly, π_Ω is linear and idempotent ($\pi_\Omega \circ \pi_\Omega = \pi_\Omega$). A tensor \mathcal{X} is in the null space of π_Ω if and only if none of its non-zero entries are indexed by N -tuples of Ω . Denoting by $\bar{\Omega}$ the complement of Ω , we have $\forall \mathcal{X}, \langle \pi_\Omega(\mathcal{X}), \pi_{\bar{\Omega}}(\mathcal{X}) \rangle = 0$.

With partial observations, we solve:

$$\min_{\mathcal{L}, \mathcal{E}} f(\mathcal{L}) + \|\mathcal{E}\|_1, \quad (23)$$

$$\text{s.t.} \quad \pi_\Omega(\mathcal{X}) = \pi_\Omega(\mathcal{L}) + \pi_\Omega(\mathcal{E}).$$

By the orthogonality of π_Ω , $\|\mathcal{E}\|_1 = \|\pi_\Omega(\mathcal{E})\|_1 + \|\pi_{\bar{\Omega}}(\mathcal{E})\|_1$. Without loss of generality we follow [7], [36] and assume $\pi_{\bar{\Omega}}(\mathcal{X}) = 0$ such that $\pi_{\bar{\Omega}}(\mathcal{X}) = \pi_{\bar{\Omega}}(\mathcal{L}) + \pi_{\bar{\Omega}}(\mathcal{E})$. This implies that we do not seek to recover possible corruption on the missing values but directly the missing element. Problem Equation (23) is therefore equivalent to:

$$\min_{\mathcal{L}, \mathcal{E}} f(\mathcal{L}) + \|\pi_\Omega(\mathcal{E})\|_1, \quad (24)$$

$$\text{s.t.} \quad \mathcal{X} = \mathcal{L} + \mathcal{E}.$$

To solve Equation (24) we need to compute the proximal operator of $\|\pi_\Omega(\mathcal{E})\|_1$, we show (proof in Appendix A.1, which can be found on the Computer Society Digital Library at <http://doi.ieeecomputersociety.org/10.1109/TPAMI.2018.2881476>) that the corresponding operator is the selective shrinkage $\mathcal{S}_\lambda^\Omega$:

$$\forall \mathcal{X} \in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}, \mathcal{S}_\lambda^\Omega(\mathcal{X}) = \mathcal{S}_\lambda(\pi_\Omega(\mathcal{X})) + \pi_{\bar{\Omega}}(\mathcal{X}). \quad (25)$$

We present tensor completion results in Section 7.4.

3.4 An Efficient Algorithm for the Degree 2 Regularizer

Finally, we derive an efficient algorithm for the degree 2-regularized problem. This discussion will serve as a basis for the other regularizers which only require minor modifications of this algorithm.

Problem (21) is not jointly convex, but is convex in each component individually. We resort to an alternating-direction method and propose a non-convex ADMM procedure that operates on the frontal slices.

Minimizing $\|\mathbf{B} \otimes \mathbf{A}\|_F$ presents a challenge: the product is high-dimensional, the two bases are coupled, and the loss is non-smooth. Let $\|\cdot\|_p$ denote the Schatten- p norm.² Using the identity $\|\mathbf{A} \otimes \mathbf{B}\|_p = \|\mathbf{A}\|_p \|\mathbf{B}\|_p$ ³ and remarking that $\|\mathbf{B}\|_F \|\mathbf{A}\|_F \leq \frac{\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2}{2}$, we minimize a simpler upper bound.⁴ The resulting sub-problems are smaller, and

2. The Schatten- p norm of \mathbf{A} is the ℓ_p norm of its singular values.

3. See Appendix E, available in the online supplemental material, and [37].

4. $\|\mathbf{A}\mathbf{B}^\top\|_*$ [38].

therefore easier to solve computationally. In order to obtain exact proximal steps for the encodings \mathbf{R}_i , we introduce a split variable \mathbf{K}_i such that $\forall i, \mathbf{K}_i = \mathbf{R}_i$. Thus, we solve:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathcal{K}, \mathcal{E}} \quad & \alpha \|\mathbf{R}\|_1 + \lambda \|\mathcal{E}\|_1 + \frac{1}{2} (\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2), \\ \text{s.t.} \quad & \mathcal{X} = \mathcal{K} \times_1 \mathbf{A} \times_2 \mathbf{B} + \mathcal{E}, \\ \text{s.t.} \quad & \mathbf{R} = \mathcal{K} \end{aligned} \quad (26)$$

By introducing the Lagrange multipliers $\mathbf{\Lambda}$ and \mathcal{Y} , such that the i th frontal slice corresponds to the i th constraint, and the dual step sizes μ and $\mu_{\mathcal{K}}$, we formulate the Augmented Lagrangian of problem Equation (26):

$$\begin{aligned} L(\mathbf{A}, \mathbf{B}, \mathbf{R}, \mathcal{E}, \mathcal{K}, \mathbf{\Lambda}, \mathcal{Y}, \mu, \mu_{\mathcal{K}}) = & \lambda \sum_i \|\mathbf{R}_i\|_1 + \lambda \sum_i \|\mathbf{E}_i\|_1 \\ & + \frac{1}{2} (\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2) + \sum_i \langle \mathbf{\Lambda}_i, \mathbf{X}_i - \mathbf{A} \mathbf{K}_i \mathbf{B}^T - \mathbf{E}_i \rangle \\ & + \sum_i \langle \mathcal{Y}_i, \mathbf{R}_i - \mathbf{K}_i \rangle + \frac{\mu}{2} \sum_i \|\mathbf{X}_i - \mathbf{A} \mathbf{K}_i \mathbf{B}^T - \mathbf{E}_i\|_F^2 \\ & + \frac{\mu_{\mathcal{K}}}{2} \sum_i \|\mathbf{R}_i - \mathbf{K}_i\|_F^2. \end{aligned} \quad (27)$$

We can now derive the ADMM updates. Each \mathbf{E}_i is given by shrinkage after rescaling [11]:

$$\mathbf{E}_i = \mathcal{S}_{\lambda/\mu}(\mathbf{X}_i - \mathbf{A} \mathbf{K}_i \mathbf{B}^T + \frac{1}{\mu} \mathbf{\Lambda}_i). \quad (28)$$

From the developments of Section 3.3, extending our algorithms to handle missing value only involves using the selective shrinkage operator Equations (25) in (28).

A similar rule is immediate to derive for \mathbf{R}_i , and solving for \mathbf{A} and \mathbf{B} is straightforward with some matrix algebra. We therefore focus on the computation of the split variable \mathbf{K}_i . Differentiating, we find \mathbf{K}_i satisfies:

$$\begin{aligned} \mu_{\mathcal{K}} \mathbf{K}_i + \mu \mathbf{A}^T \mathbf{A} \mathbf{K}_i \mathbf{B}^T \mathbf{B} \\ = \mathbf{A}^T (\mathbf{\Lambda}_i + \mu (\mathbf{X}_i - \mathbf{E}_i)) \mathbf{B} + \mu_{\mathcal{K}} \mathbf{R}_i + \mathcal{Y}_i. \end{aligned} \quad (29)$$

The key is here to recognize Equation (29) is a *Stein* equation, and can be solved in cubical time and quadratic space in r by solvers for discrete-time Sylvester equations - such as the Hessenberg-Schur method [39] - instead of the naive $O(r^6)$ time, $O(r^4)$ space solution of vectorizing the equation in a size- r^2 linear system. We obtain Algorithm 1.

4 RKCA AND GLOBAL OPTIMALITY

The work of [30], [31] suggests global optimality can be achieved from any initialization in tensor factorization models given that the factorization and regularization mappings match in certain ways. We summarize the main results of this work and show our model with regularizer Equation (15) or (16) respects the conditions for global optimality.

4.1 Review of the Main Results

We first review some of the concepts manipulated in [30], [31] and give an overview of the main results.

Algorithm 1. RKCA with Degree 2 Regularization

```

1: procedure RKCA( $\mathcal{X}, r, \lambda, \alpha$ )
2:    $\mathbf{A}^0, \mathbf{B}^0, \mathcal{E}^0, \mathcal{R}^0, \mathcal{K}^0, \mathbf{\Lambda}^0, \mathcal{Y}^0, \mu^0, \mu_{\mathcal{K}}^0 \leftarrow \text{INITIALIZE}(\mathcal{X})$ 
3:   while not converged do
4:      $\mathcal{E}^{t+1} \leftarrow \mathcal{S}_{\lambda/\mu^t}(\mathcal{X} - \mathcal{K}^t \times_1 \mathbf{A}^t \times_2 \mathbf{B}^t + \frac{1}{\mu^t} \mathbf{\Lambda}^t)$ 
5:      $\tilde{\mathcal{X}}^{t+1} \leftarrow \mathcal{X} - \mathcal{E}^{t+1}$ 
6:      $\mathbf{A}^{t+1} \leftarrow (\sum_i (\mu^t \tilde{\mathbf{X}}_i^{t+1} + \mathbf{\Lambda}_i^t) \mathbf{B}^t (\mathbf{K}_i^t)^T) /$ 
        $(\mathbf{I} + \mu^t \sum_i \mathbf{K}_i^t (\mathbf{B}^t)^T \mathbf{B}^t (\mathbf{K}_i^t)^T)$ 
7:      $\mathbf{B}^{t+1} \leftarrow (\sum_i (\mu^t \tilde{\mathbf{X}}_i^{t+1} + \mathbf{\Lambda}_i^t)^T \mathbf{A}^{t+1} \mathbf{K}_i^t) /$ 
        $(\mathbf{I} + \mu^t \sum_i (\mathbf{K}_i^t)^T (\mathbf{A}^{t+1})^T \mathbf{A}^{t+1} \mathbf{K}_i^t)$ 
8:     for all  $i$  do
9:        $\mathbf{K}_i^{t+1} \leftarrow \text{STEIN}(-\frac{\mu^t}{\mu_{\mathcal{K}}} (\mathbf{A}^{t+1})^T \mathbf{A}^{t+1}, (\mathbf{B}^{t+1})^T \mathbf{B}^{t+1},$ 
          $\frac{1}{\mu_{\mathcal{K}}} [(\mathbf{A}^{t+1})^T (\mathbf{\Lambda}_i^t + \mu^t \tilde{\mathbf{X}}_i^{t+1}) \mathbf{B}^{t+1} + \mathcal{Y}_i^t] + \mathbf{R}_i^t)$ 
10:       $\mathbf{R}_i^{t+1} \leftarrow \mathcal{S}_{\alpha/\mu_{\mathcal{K}}}^t(\mathbf{K}_i^{t+1} - \frac{1}{\mu_{\mathcal{K}}} \mathcal{Y}_i^t)$ 
11:     end for
12:      $\mathbf{\Lambda}^{t+1} \leftarrow \mathbf{\Lambda}^t + \mu^t (\tilde{\mathcal{X}}^{t+1} - \mathcal{K}^{t+1} \times_1 \mathbf{A}^{t+1} \times_2 \mathbf{B}^{t+1})$ 
13:      $\mathcal{Y}^{t+1} \leftarrow \mathcal{Y}^t + \mu_{\mathcal{K}}^t (\mathcal{R}^{t+1} - \mathcal{K}^{t+1})$ 
14:      $\mu^{t+1} \leftarrow \min(\mu^*, \rho \mu^t)$ 
15:      $\mu_{\mathcal{K}}^{t+1} \leftarrow \min(\mu_{\mathcal{K}}^*, \rho \mu_{\mathcal{K}}^t)$ 
16:   end while
17:   return  $\mathbf{A}, \mathbf{B}, \mathcal{R}, \mathcal{E}$ 
18: end procedure

```

Definition 4.1 (Def. 1 of [31, page 6]). A size- r set of K factors $(\mathcal{X}^1, \dots, \mathcal{X}^K)_r$ is defined to be a set of K tensors where the final dimension of each tensor is equal to r : $(\mathcal{X}^1, \dots, \mathcal{X}^K)_r \in \mathbb{R}^{(D^1 \times r)} \times \dots \times \mathbb{R}^{(D^K \times r)}$.

Definition 4.2 (Def. 6 of [31, page 6]). An elemental mapping, $\phi: \mathbb{R}^{D^1} \times \dots \times \mathbb{R}^{D^K} \rightarrow \mathbb{R}^D$ is any mapping which is positively homogeneous with degree $p \neq 0$. The r -element factorization mapping $\Phi_r: \mathbb{R}^{(D^1 \times r)} \times \dots \times \mathbb{R}^{(D^K \times r)} \rightarrow \mathbb{R}^D$ is defined as:

$$\Phi_r(\mathcal{X}^1, \dots, \mathcal{X}^K) = \sum_{i=1}^r \phi(\mathcal{X}_i^1, \dots, \mathcal{X}_i^K). \quad (30)$$

Definition 4.3 (Def. 7 of [31, page 8]). An elemental regularization function $g: \mathbb{R}^{D^1} \times \dots \times \mathbb{R}^{D^K} \rightarrow \mathbb{R}_+ \cup \infty$, is defined to be any function which is positive semidefinite and positively homogeneous.

Definition 4.4 (Def. 8 of [31, page 9]). Given an elemental mapping ϕ and an elemental regularization function g , the authors define (ϕ, g) to be a nondegenerate pair if 1) g and ϕ are both positively homogeneous with degree p , for some $p \neq 0$ and 2) $\forall X \in \text{Im}(\phi) \setminus 0, \exists \mu \in (0, \infty]$ and $(\tilde{z}^1, \dots, \tilde{z}^K)$ such that $\phi(\tilde{z}^1, \dots, \tilde{z}^K) = X, g(\tilde{z}^1, \dots, \tilde{z}^K) = \mu$, and $g(z^1, \dots, z^K) \geq \mu$ for all (z^1, \dots, z^K) such that $\phi(z^1, \dots, z^K) = X$.

The main result, Theorem 15 of [31], provides a characterization of the global optima of CP-based tensor factorization problems.

Theorem 4.5 (Theorem 15 of [31, page 15]). Given a function $f_r(\mathcal{X}^1, \dots, \mathcal{X}^K, Q)$, any local minimizer of the optimization problem:

$$\begin{aligned} \min_{(\mathcal{X}^1, \dots, \mathcal{X}^K)_r, Q} \quad & f_r(\mathcal{X}^1, \dots, \mathcal{X}^K, Q) \equiv \\ \ell(\Phi_r(\mathcal{X}^1, \dots, \mathcal{X}^K), Q) + \lambda \sum_{i=1}^r g(\mathcal{X}_i^1, \dots, \mathcal{X}_i^K) + H(Q) \end{aligned} \quad (31)$$

such that $(\mathcal{X}_{i_0}^1, \dots, \mathcal{X}_{i_0}^K) = (0, \dots, 0)$ for some $i_0 \in \{1, \dots, r\}$ is a global minimizer.

Where Q is an optional set of non-factorized variables (in our case \mathcal{E}), H is convex and possibly non-smooth, ℓ is jointly convex and once-differentiable in (\mathcal{X}, Q) , and (ϕ, g) is a nondegenerate pair.

4.2 Outline of the Proof

In order to apply the aforementioned results, we reformulate RKCA in an equivalent formulation that satisfies the hypotheses of [30], [31]. The arguments we develop actually hold for the general Tucker factorization. This section outlines the proof of equivalence, the detailed developments can be found in Appendix C, available in the online supplemental material.

The main challenge is to find a factorization function that can be expressed as a sum of elemental mappings. Problem (21) as-is does not lend itself to such a formulation: even though the factors \mathbf{A} and \mathbf{B} form a size- r set of 2 factors, the core tensor \mathcal{R} is of size N on its last dimension. We note however that the set of frontal slices of \mathcal{R} is a size- r set of N factors, but this formulation doesn't satisfy Proposition 10 of [31] and it is not immediately obvious how to define concatenation of the factors of \mathcal{X} and \mathcal{Y} and how to verify the convexity of the factorization-regularization function Ω (Equation (18) and Proposition 11 of [31] omitted here for brevity). Additionally, the results of [30], [31] have been proved in the case of a single sum over the last dimension of the factors, but our factorization being a special case of Tucker decomposition is most naturally described as:

$$\mathcal{L} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} \quad (32)$$

$$= \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{I}_N \quad (33)$$

$$= \sum_i \sum_j \sum_k r_{i,j,k} \cdot \mathbf{a}_i \otimes \mathbf{b}_j \otimes \mathbf{c}_k. \quad (34)$$

First, we note the order of the sums can be permuted such that it is clear our model expresses \mathcal{L} as the sum of N tensors of size $m \times n \times N$, where all frontal slices are the null matrix except for one:

$$\mathcal{L} = \sum_k \sum_i \sum_j r_{i,j,k} \cdot \mathbf{a}_i \otimes \mathbf{b}_j \otimes \mathbf{c}_k. \quad (35)$$

With $(\mathbf{c}_k)_i = \delta_{i,k} = \begin{cases} 1 & i = k \\ 0 & i \neq k \end{cases}$.

Next, we seek a transformation such that Equation (35) does not involve cross-terms. Our idea is to unfold the Tucker factorization by duplicating elements of the factors to express it in the form Equation (30) where the sum is over a size- s set of 3 factors (or M factors for the general Tucker).

- We define $\sigma = \text{vec}(\mathcal{R})$.
- We construct $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ from \mathbf{A} and \mathbf{B} by duplicating columns.
- \mathbf{I}_N doesn't depend on the data and can be injected directly in the elemental mapping.

We show the factorization function is defined over the size- Nr^2 set of 3 factors $(\sigma, \tilde{\mathbf{A}}, \tilde{\mathbf{B}})$ by:

$$\Phi_{Nr^2}(\sigma, \tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \sum_{l=1}^{Nr^2} \sigma_l \cdot \tilde{\mathbf{a}}_l \otimes \tilde{\mathbf{b}}_l \otimes \delta_{l, (\lceil \frac{l}{Nr} \rceil - 1) \bmod r^2 + 1}. \quad (36)$$

4.3 RKCA with Degree 3 Regularization

We now check that the *degree 3 regularizers* introduced in Section 3 are compatible with our reformulated factorization. Recall both regularizers are of the form:

$$g(\mathcal{L}) = \alpha \|\mathcal{R}\|_1 \|\mathbf{A}\|_q \|\mathbf{B}\|_q. \quad (37)$$

With $\|\cdot\|_p$ the Schatten- p norm, $p = 1$ for the Nuclear norm and $p = 2$ for the Frobenius norm. As we shall see in Section 5.2, g is low-rank promoting.

In the case of the Frobenius penalty, the resulting optimization problem is:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}, \mathcal{E}} \quad & \lambda \|\mathcal{R}\|_1 \|\mathbf{B}\|_F \|\mathbf{A}\|_F + \lambda \|\mathcal{E}\|_1, \\ \text{s.t.} \quad & \mathcal{X} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} + \mathcal{E}. \end{aligned} \quad (38)$$

The regularizer g adapted to our transformation is:

$$g(\sigma, \tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \|\sigma\|_1 \frac{\|\tilde{\mathbf{A}}\|_F}{Nr} \frac{\|\tilde{\mathbf{B}}\|_F}{Nr}. \quad (39)$$

This stems from the fact that the Frobenius norm is equivalent to the element-wise ℓ_2 norm, and each element of \mathbf{A} and \mathbf{B} appears Nr times in $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$.

In the case of the Nuclear norm, we simply observe that the numbers of linearly independent columns in $\tilde{\mathbf{A}}$ and $\tilde{\mathbf{B}}$ are clearly the same as in \mathbf{A} and \mathbf{B} , so our transformation preserves the ranks of the target matrices, and the regularization function is simply:

$$f(\sigma, \tilde{\mathbf{A}}, \tilde{\mathbf{B}}) = \|\sigma\|_1 \|\tilde{\mathbf{A}}\|_* \|\tilde{\mathbf{B}}\|_*. \quad (40)$$

It should be clear that the newly defined ϕ and g are both positively homogeneous of degree 3, and form a nondegenerate pair. Property 10 of [31] also holds.

Hence, we argue that RKCA with a product of norms regularization enjoys the optimality guarantees of [31].

5 RKCA LEARNS LOW-RANK DICTIONARIES

In this Section we explain the low-rank promoting behavior of RKCA from complementary perspectives. First, we show the regularizers defined in Equations (14), (15), and (16) directly provide upper bounds on the mode-1 and mode-2 ranks of the low-rank component, and thus on the rank of each of its frontal slices. This perspective was first presented in [29]. The second approach to explaining the models' properties studies the optimization sub-problems associated with the two bases \mathbf{A} and \mathbf{B} . Based on recent work [40], we show these sub-problems are equivalent to rank-minimization problems and admit closed-form solutions that involve forms of singular value thresholding. Finally, we discuss connections with recent work on low-rank inducing norms.

5.1 Direct Penalization of the Rank of \mathcal{L}

Seeing the models from the perspective of Robust PCA, which seeks a low-rank representation \mathbf{A} of the dataset \mathbf{X} , we minimize the rank of the low-rank tensor \mathcal{L} . More precisely, we show in Theorem 5.1 that we simultaneously penalize the Tucker rank and the multi-rank of \mathcal{L} .

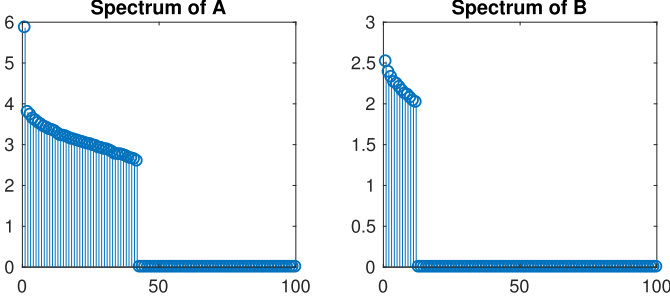


Fig. 3. Sample spectrums of \mathbf{A} and \mathbf{B} . Ground truth attained (\mathbf{A} : 42, \mathbf{B} : 12). $r = 100$. Degree 2 regularizer.

Theorem 5.1. *RKCA encourages low mode-1 and mode-2 rank, and thus, low-rankness in each frontal slice of \mathcal{L} , for suitable choices of the parameters λ and α .*

Proof. We minimize either $\lambda\|\mathcal{E}\|_1 + \alpha\|\mathcal{R}\|_1 + \|\mathbf{A} \otimes \mathbf{B}\|_F$, $\lambda\|\mathcal{E}\|_1 + \alpha\|\mathcal{R}\|_1 + \|\mathbf{A} \otimes \mathbf{B}\|_F$, or $\lambda\|\mathcal{E}\|_1 + \alpha\|\mathcal{R}\|_1 + \|\mathbf{A} \otimes \mathbf{B}\|_*$. From the equivalence of norms in finite-dimensions, $\exists k \in \mathbb{R}_+$, $\|\mathbf{A} \otimes \mathbf{B}\|_* \leq k\|\mathbf{A} \otimes \mathbf{B}\|_F$. In the case of the Frobenius norm, we can choose $\alpha = \frac{\alpha'}{k}$, $\lambda = \frac{\lambda'}{k}$ to reduce the problem to that of the nuclear norm. In all cases, we penalize $\text{rank}(\mathbf{A} \otimes \mathbf{B}) = \text{rank}(\mathbf{A})\text{rank}(\mathbf{B})$. Given that the rank is a non-negative integer, $\text{rank}(\mathbf{A})$ or $\text{rank}(\mathbf{B})$ decreases necessarily. Therefore, we minimize the mode-1 and mode-2 ranks of $\mathcal{L} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B}$. Additionally, $\forall i, \text{rank}(\mathbf{A}\mathbf{R}_i\mathbf{B}^\top) \leq \min(\text{rank}(\mathbf{A}), \text{rank}(\mathbf{B}), \text{rank}(\mathbf{R}_i))$.

Exhibiting a valid k may help in the choice of parameters. We show $\forall \mathbf{A} \in \mathbb{R}^{m \times n}$, $\|\mathbf{A}\|_* \leq \sqrt{\min(m, n)}\|\mathbf{A}\|_F$. We know $\forall \mathbf{x} \in \mathbb{R}^n$, $\|\mathbf{x}\|_1 \leq \sqrt{n}\|\mathbf{x}\|_2$. Recalling the nuclear norm and the Frobenius norm are the Schatten-1 and Schatten-2 norms, and \mathbf{A} has $\min(m, n)$ singular values, the result follows. \square

In practice, we find that on synthetic data designed to test the models, we effectively recover the ranks of \mathbf{A} and \mathbf{B} regardless of the choice of r , as seen in Fig. 3.

5.2 Analysis of the Sub-Problems in \mathbf{A} and \mathbf{B}

We present an alternative interpretation of the low-rank properties of our methods. In [40] the authors show that Frobenius norm and Nuclear norm regularizations either lead to the same optimal solution - even in the presence of noise - when the dictionary provides enough representative power, or are equivalent in the sense that they describe two different bases on the column space of the dictionary. We explain the behavior of our algorithms by showing the sub-problems in \mathbf{A} and \mathbf{B} are rank-minimization problems.

5.2.1 Sub-Problems in \mathbf{A} and \mathbf{B}

For generality purposes, let us begin with the problem:

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}, \mathcal{R}, \mathcal{E}} \quad & \beta(\mathcal{R})\|\mathbf{A}\|_F\|\mathbf{B}\|_F + \lambda\|\mathcal{E}\|_1 + W(\mathcal{R}), \\ \text{s.t.} \quad & \mathcal{X} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{I}_N + \mathcal{E}. \end{aligned} \quad (41)$$

In the degree 3 case we have $\beta(\mathcal{R}) = \alpha\|\mathcal{R}\|_1$ and $W(\mathcal{R}) = 0$, in the degree 2 case, $\beta(\mathcal{R}) = 1$ and $W(\mathcal{R}) = \alpha\|\mathcal{R}\|_1$. We now omit the dependency in \mathcal{R} for β and W as they are constant in the sub-problems in \mathbf{A} and \mathbf{B} .

The sub-problem in \mathbf{A} is:

$$\begin{aligned} \min_{\mathbf{A}} \quad & [\beta\|\mathbf{B}\|_F]\|\mathbf{A}\|_F, \\ \text{s.t.} \quad & \mathcal{X} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{I}_N + \mathcal{E}. \end{aligned} \quad (42)$$

Equivalently, in a matricized way using Property (2.1):

$$\begin{aligned} \min_{\mathbf{A}} \quad & [\beta\|\mathbf{B}\|_F]\|\mathbf{A}\|_F, \\ \text{s.t.} \quad & \mathbf{X}_{[1]} = \mathbf{A}\mathbf{R}_{[1]}(\mathbf{I}_N \otimes \mathbf{B})^\top + \mathbf{E}_{[1]}. \end{aligned} \quad (43)$$

Noting that $\|\mathbf{A}\|_F = \|\mathbf{A}^\top\|_F$ we reformulate the problem in the form of Equation (4) of [40]:

$$\min_{\mathbf{A}} \gamma\|\mathbf{A}^\top\|_F, \quad \text{s.t.} \quad \tilde{\mathbf{X}}_{[1]}^\top = \mathbf{D}_\mathbf{A}\mathbf{A}^\top. \quad (44)$$

With $\tilde{\mathbf{X}}_{[1]}^\top = \mathbf{X}_{[1]}^\top - \mathbf{E}_{[1]}^\top$, $\gamma = [\beta\|\mathbf{B}\|_F]$, $\mathbf{D}_\mathbf{A} = (\mathbf{I}_N \otimes \mathbf{B})\mathbf{R}_{[1]}^\top$.

The sub-problem in \mathbf{B} is very similar:

$$\begin{aligned} \min_{\mathbf{B}} \quad & [\beta\|\mathbf{A}\|_F]\|\mathbf{B}\|_F, \\ \text{s.t.} \quad & \mathcal{X} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} \times_3 \mathbf{I}_N + \mathcal{E}. \end{aligned} \quad (45)$$

Matricizing on the second mode:

$$\min_{\mathbf{B}} \kappa\|\mathbf{B}^\top\|_F, \quad \text{s.t.} \quad \tilde{\mathbf{X}}_{[2]}^\top = \mathbf{D}_\mathbf{B}\mathbf{B}^\top. \quad (46)$$

With $\tilde{\mathbf{X}}_{[2]}^\top = \mathbf{X}_{[2]}^\top - \mathbf{E}_{[2]}^\top$, $\kappa = [\beta\|\mathbf{A}\|_F]$, $\mathbf{D}_\mathbf{B} = (\mathbf{I}_N \otimes \mathbf{A})\mathbf{R}_{[2]}^\top$.

5.2.2 Optimal Solutions and Interpretation

According to [17] the optimal solution of $\min\|\mathbf{C}\|_*$, s.t. $\mathbf{X} = \mathbf{D}\mathbf{C}$ assuming feasible solutions exist and $\mathbf{D} \neq \mathbf{0}$ is $\mathbf{C}^* = \mathbf{V}_r\mathbf{V}_r^\top$ where $\mathbf{U}_r\mathbf{\Sigma}_r\mathbf{V}_r^\top$ is the skinny SVD of \mathbf{D} . [41] showed it is also the optimal solution of $\min\|\mathbf{C}\|_F$, s.t. $\mathbf{X} = \mathbf{D}\mathbf{C}$.

It is therefore easy to show that the optimal solution of $\min\|\mathbf{C}^\top\|_*$, s.t. $\mathbf{X} = \mathbf{D}\mathbf{C}^\top$ under similar conditions is $\mathbf{U}_r\mathbf{U}_r^\top$.

Proof. Letting $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ we have $\mathbf{X}^\top = \mathbf{V}\mathbf{\Sigma}\mathbf{U}^\top$. \square

Letting $\mathbf{U}_p^\mathbf{A}\mathbf{\Sigma}_p^\mathbf{A}\mathbf{V}_p^\mathbf{A}^\top$ the skinny SVD of $\mathbf{D}_\mathbf{A}$ and $\mathbf{U}_q^\mathbf{B}\mathbf{\Sigma}_q^\mathbf{B}\mathbf{V}_q^\mathbf{B}^\top$ the skinny SVD of $\mathbf{D}_\mathbf{B}$ the optimal solutions for \mathbf{A} and \mathbf{B} are therefore $\mathbf{U}_p^\mathbf{A}\mathbf{U}_p^\mathbf{A}^\top$ and $\mathbf{U}_q^\mathbf{B}\mathbf{U}_q^\mathbf{B}^\top$.

For the sake of brevity, we shall describe the interpretation for \mathbf{A} as the same holds for \mathbf{B} by symmetry. It should be noted that $\mathbf{U}_p^\mathbf{A}$ is the PCA basis of $\mathbf{D}_\mathbf{A}$ and that $\mathbf{U}_p^\mathbf{A}\mathbf{U}_p^\mathbf{A}^\top$ is the matrix of the orthogonal projection onto that basis. Thus, the optimal \mathbf{A} is the projection of $\mathbf{D}_\mathbf{A}$ onto its principal components. Remembering that $\mathbf{D}_\mathbf{A} = (\mathbf{I}_N \otimes \mathbf{B})\mathbf{R}_{[1]}^\top$, $\mathbf{D}_\mathbf{A}$ is the product of \mathbf{B} with each code \mathbf{R}_i matricized such that the i th line of $\mathbf{D}_\mathbf{A}$ is the concatenation of all the i th columns of all the partial reconstructions $\mathbf{R}_k\mathbf{B}^\top$, and all the columns of $\mathbf{D}_\mathbf{A}$ are all the lines of $\mathbf{R}_k\mathbf{B}^\top$, which can be seen as the partial reconstruction of the low-rank images in their common row space $\text{Span}(\mathbf{B})$.

Hence, we can see the process of updating \mathbf{A} and \mathbf{B} as alternating PCAs in the row and column spaces, respectively.

5.3 A Note on Low-Rank Inducing Norms

The recently published [42] suggests one more interpretation of our regularizers. In [42], the authors show a family of low-rank inducing norms, based on solving rank-constrained

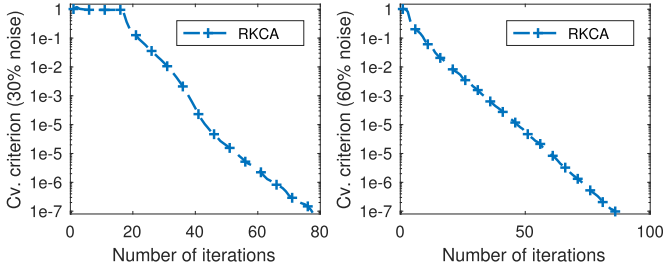


Fig. 4. Convergence on synthetic data with 30 and 60 percent corruption.

Frobenius norm minimization, or rank-constrained Spectral norm minimization problems can outperform the standard nuclear norm relaxation in rank-minimization problems. Given that the role of the parameter r in our model is to impose an upper bound on the ranks of the dictionaries and of the reconstructions, we suggest the choice of Frobenius regularization instead of nuclear norm regularization is further justified. Deeper analysis of this connection is required and is left for future work.

6 IMPLEMENTATION DETAILS AND COMPLEXITY

In this section, we discuss the computational complexity of our algorithms, and the scalability issues depending on the regularization. We discuss parameter tuning and adaptive ADMM schemes [43], [44], [45] in Appendix F.1, available in the online supplemental material. The initialization procedure has already been discussed in [29], and is included in Appendix F.2, available in the online supplemental material along with references regarding the convergence of ADMM solvers [46], [47]. Empirically, we obtained systematic convergence of Algorithm 1 to a good solution, and a linear convergence rate, as shown in Fig. 4.

6.1 Computational Complexity and LADMM

The substitution method used in Section 3.4 is effective but comes with an additional cost that limits its scalability. Notably, the cost of solving N Stein equations cannot be neglected in practice. Additionally, the added parameters can make tuning difficult. Linearization provides an alternative approach that can scale better to large dimensions and large datasets. In Appendix D, available in the online supplemental material, we give detailed derivations of how RKCA can be implemented with LADMM.

The time and space complexity per iteration of Algorithm 1 (i.e., with the degree two regularizer) are $O(N(mnr + (m + n)r + mn + \min(m, n)r^2 + r^3 + r^2))$ and $O(N(mn + r^2) + (m + n)r + r^2)$. Since $r \leq \min(m, n)$, the terms in r are asymptotically negligible, but in practice it is useful to know how the computational requirements scale with the size of the dictionary. Similarly, the initialization procedure has cost $O(N(mn \min(m, n) + (\min(m, n))^3 + mn) + mn)$ in time and needs quadratic space per slice, assuming a standard algorithm is used for the SVD[48].

Switching to an LADMM update for \mathcal{R} eliminates the need of solving costly Stein equations. The soft-shrinkage operator is applied to the tensor $(\mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B} - \Delta) \times_1 \mathbf{A}^T \times_2 \mathbf{B}^T$, which has Nr^2 elements, and can be computed in $O(N(\min(m, n)r^2 + mn + r^2 + r^3) + mr^2 + nr^2)$ by remembering that $(\mathcal{X} \times_1 \mathbf{A})$

$\times_1 \mathbf{B} = \mathcal{X} \times_1 \mathbf{B} \mathbf{A}$. The space complexity of the update is $O(N(\min(m, n)r + r^2) + r^2)$.

Updating \mathbf{A} and \mathbf{B} with a Frobenius norm or with a nuclear norm penalty requires computing a proximal operator, and solving a linear system in the case of ADMM with substitution. We focus on the computation of the proximal operator in Section 6.2. The substitution adds an additional time and space complexity of $O(mr + nr)$.

Several key steps of Algorithm 1 and its variants, such as summations of independent terms, are trivially distributed in a *MapReduce* [49] way. Proximal operators are separable in nature and are therefore parallelizable. Consequently, highly parallel and distributed implementations are possible, and computational complexity can be further reduced by adaptively adopting sparse linear algebra structures.

6.2 Frobenius and Nuclear Norm Penalties

In Equation (26) we used an upper bound on $\|\mathbf{A}\|_F \|\mathbf{B}\|_F$ to obtain smooth sub-problems. In the degree 3-regularized case, we did not apply the same bound. Given the non-smoothness of the Frobenius norm, we must resort to other approaches such as substitution or linearization (c.f. Appendix B and Appendix D, available in the online supplemental material), as with the Nuclear norm.

In Section 5.2, we then showed how Frobenius and Nuclear norm penalties can yield the same optimal solutions, and how in our case they would be equivalent in solving the sub-problems associated with \mathbf{A} and \mathbf{B} . It is therefore natural to wonder why we should choose one over the other, and what the practical implications of this choice are.

It can be shown that the proximal operator of the Schatten- p norm has a closed form expression that requires the computation of the SVD of the matrix (Proposition A.1, Appendix A.2, available in the online supplemental material). Computing the SVD is costly, with a typical algorithm [48] requiring $O(mn \min(m, n) + (\min(m, n))^3)$ floating-point operations. Storing the resulting triple $(\mathbf{U}, \mathbf{S}, \mathbf{V})$ requires $m \times \min(m, n) + \min(m, n)^2 + n \times \min(m, n)$ space. Although faster algorithms have been developed, scalability remains a concern.

However, the Frobenius norm is equal to the matrix element-wise ℓ_2 norm, whose proximal operator is only the projection on the unit ball and doesn't require any costly matrix decomposition. The choice of the Frobenius norm is therefore justified in the high dimensional setting. A comparable use of the Frobenius norm in lieu of the Nuclear norm for representation learning can be seen in [50], also motivated by scalability.

7 EXPERIMENTAL EVALUATION

We first provide experimental verification of the correctness of Algorithm 1 (with degree 2 regularization) on synthetic data. We then compared the performance of RKCA with degree 2 regularization against a range of state-of-the-art tensor decomposition algorithms on four low-rank modeling computer vision benchmarks: two for image denoising, and two for background subtraction. As a baseline, we report the performance of matrix Robust PCA implemented via inexact ALM (RPCA) [11], [51], of Non-Negative Robust Dictionary Learning (RNNDL) [52], and of the Low-Rank Representations algorithm [17] implemented with both exact ALM

(*LRR*) and inexact ALM (*LRRi*). We chose the following methods to include recent representatives of various existing approaches to low-rank modeling on tensors: The singleton version of Higher-Order Robust PCA (*HORPCA-S*) [7] optimizes the Tucker rank of the tensor through the sum of the nuclear norms of its unfoldings. In [53], the authors consider a similar model but with robust M-estimators as loss functions, either a Cauchy loss or a Welsh loss, and support both hard and soft thresholding; we tested the soft-thresholding models (*Cauchy ST* and *Welsh ST*). Non-convex Tensor Robust PCA (*NC TRPCA*) [54] adapts to tensors the matrix non-convex RPCA [55]. Finally, the two Tensor RPCA algorithms [34], [35] (*TRPCA '14* and *TRPCA '16*) work with slightly different definitions of the tensor nuclear norm as a convex surrogate of the tensor multi-rank. In addition to the aforementioned low-rank matrix and tensor factorizations, we include in the comparison the recently proposed Deep Image Prior [32] with two different architectures (*DIP 1* and *DIP 2*) as an example of deep learning based approach for denoising. Our choice of comparing to DIP is motivated by the fact that DIP is fully unsupervised, designed for denoising (among other tasks), does not require a large dataset for training, and is close in spirit to tensor factorizations. Given the small sizes of the datasets at hand, training deep neural networks would have been impractical.

For all models but RNNDL, we used the implementation made available by their respective authors, either publicly or on request. Our implementation of RNNDL was tested for conformity with the original paper.

For each model, with the exception of DIP, we identified a maximum of two parameters to tune via grid-search in order to keep parameter tuning tractable. When criteria or heuristics for choosing the parameters were provided by the authors, we chose the search space around the value obtained from them. In all cases, the tuning process explored a wide range of parameters to maximize performance. In the special case of DIP, we re-used the two architectures implemented in the denoising example in the code made available online by the authors (c.f. Appendix G, available in the online supplemental material). We then trained for 2500 iterations while keeping track of the best reconstruction so-far measured by the PSNR. The best reconstruction is then saved for further comparison. Following [32], the input is chosen to be random. As in [17], the data matrix itself was used as the dictionary for LRR and LRRi.

When the performance of one method was significantly worse than that of the other, the result is not reported so as not to clutter the text (see Appendix H, available in the online supplemental material). This is the case of Separable Dictionary Learning [18] whose drastically different nature renders unsuitable for robust low-rank modeling, but was compared for completeness. For the same reason, we did not compare our method against K-SVD [13], or [19].

Finally, we provide tensor completion experiments with and without gross corruption in Section 7.4.

7.1 Validation on Synthetic Data

We generated synthetic data following the RKCA model's assumptions by first sampling two random bases \mathbf{A} and \mathbf{B} of known ranks r_A and r_B , N Gaussian slices for the core \mathcal{R} , and forming the ground truth $\mathcal{L} = \mathcal{R} \times_1 \mathbf{A} \times_2 \mathbf{B}$. We modeled

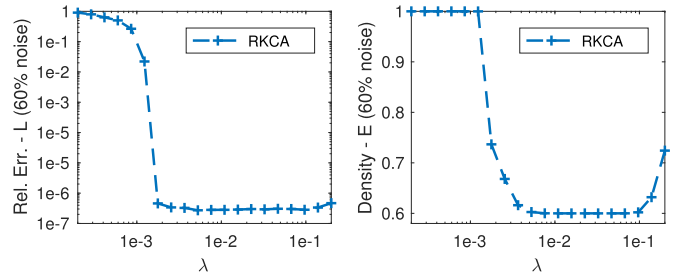


Fig. 5. Recovery performance with 60 percent corruption. Relative ℓ_2 error and density.

additive random sparse Laplacian noise with a tensor \mathcal{E} whose entries are 0 with probability p , and 1 or -1 with equal probability otherwise. We generated data for $p = 70\%$ and $p = 40\%$, leading to a noise density of, respectively, 30 and 60 percent. We measured the reconstruction error on \mathcal{L} and \mathcal{E} , and the density of \mathcal{E} for varying values of λ , and $\alpha = 10^{-2}$. Our model achieved near-exact recovery of both \mathcal{L} and \mathcal{E} , and exact recovery of the density of \mathcal{E} , for suitable values of λ . Evidence is presented in Fig. 5 for the 60 percent noise case.

Algorithm 1 appears robust to small changes in λ , which suggests more than one value can lead to optimal results, and that a simple criterion that provides consistently good reconstruction might be derived, as in Robust PCA [11]. In the 30 percent noise case, we did not observe an increase in the density of \mathcal{E} as λ increases, and the ℓ_2 error on both \mathcal{E} and \mathcal{L} was of the order of 1×10^{-7} .

7.2 Background Subtraction

Background subtraction is a common task in computer vision and can be tackled by robust low-rank modeling: the static or mostly static background of a video sequence can effectively be represented as a low-rank tensor while the foreground forms a sparse component of outliers.

7.2.1 Experimental Procedure

We compared the algorithms on two benchmarks. The first is an excerpt of the *Highway* dataset [56], and consists in a video sequence of cars travelling on a highway; the background is completely static. We kept 400 gray-scale images re-sized to 48×64 pixels. The second is the *Airport Hall* dataset ([57]) and has been chosen as a more challenging benchmark since the background is not fully static and the scene is richer. We used the same excerpt of 300 frames (frames 3301 to 3600) as in [58], and kept the frames in their original size of 144×176 pixels.

We treat background subtraction as a binary classification problem. Since ground truth frames are available for our excerpts, we report the AUC [59] on both videos. The value of α was set to 1×10^{-2} for both experiments.

7.2.2 Results

The original, ground truth, and recovered frames are in Fig. 6 for the *Hall* experiment (*Highway* in Appendix H, available in the online supplemental material).

Table 1 presents the AUC scores of the algorithms, ranked in order of their mean performance on the two benchmarks. The two matrix methods rank high on both benchmarks and only half of the tensor algorithms match or outperform this baseline. Our proposed model matches the best performance

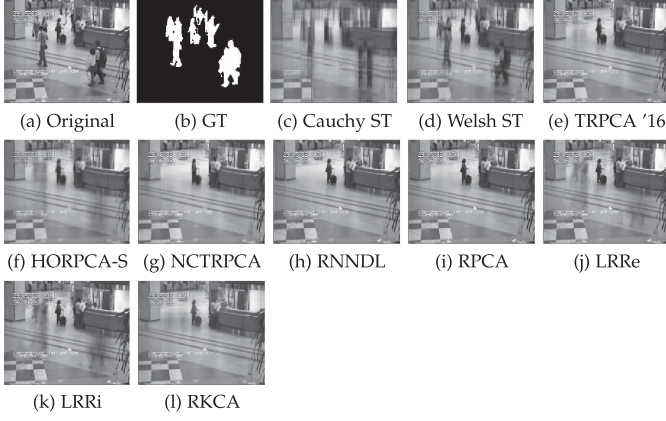


Fig. 6. Background subtraction results on *Airport Hall*. TRPCA '14 removed.

on the *Highway* dataset and provides significantly higher performance than the other on the more challenging *Hall* benchmark. Visual inspection of the results show RKCA is the only method that doesn't fully capture the immobile people in the background, and therefore achieves the best trade-off between foreground detection and background-foreground contamination.

7.3 Image Denoising

Many natural and artificial images exhibit an inherent low-rank structure and are suitably denoised by low-rank modeling algorithms. In this section, we assess the performance of the cohort on two datasets chosen for their popularity, and for the typical use cases they represent.

We consider collections of grayscale images, and color images represented as 3-way tensors. Laplacian (salt & pepper) noise was introduced separately in all frontal slices of the observation tensor at three different levels: 10, 30, and 60 percent, to simulate medium, high, and gross corruption. In these experiments we set the value of α to 1×10^{-3} for noise levels up to 30 percent, and to 1×10^{-2} at the 60 percent level.

We report two complementary image quality metrics. The *Peak Signal To Noise Ratio (PSNR)* will be used as an indicator of the element-wise reconstruction quality of the signals, while the *Feature Similarity Index (FSIM, FSIMc for color images)* [60] evaluates the recovery of structural information. Quantitative metrics are not perfect replacements for subjective assessment of image quality; therefore, we present reconstructed images for verification. Our measure of choice

TABLE 1
AUC on *Highway* and *Hall* Ordered by Mean AUC

Algorithm	Highway	Hall
RKCA (proposed)	0.94	0.88
TRPCA '16	0.94	0.86
NC TRPCA	0.93	0.86
RPCA (baseline)	0.94	0.85
RNNDL (baseline)	0.94	0.85
LRR Exact (baseline)	0.94	0.84
LRR Inexact (baseline)	0.93	0.84
HORPCA-S	0.93	0.86
Cauchy ST	0.83	0.76
Welsh ST	0.82	0.71
TRPCA '14	0.76	0.61

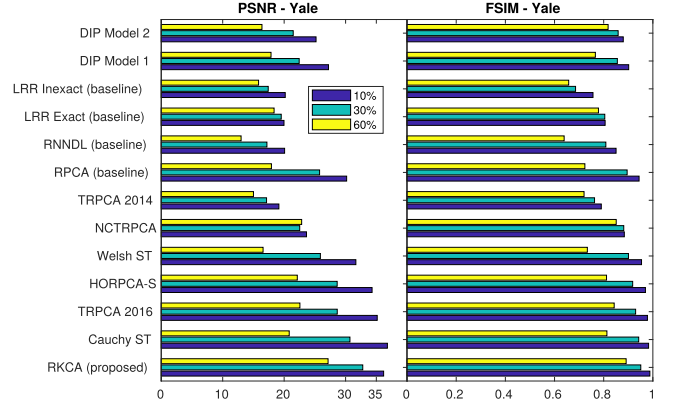


Fig. 7. Mean PSNR and FSIM on the 64 images of the first subject of *Yale* at noise levels 10, 30, and 60 percent.

for determining which images to compare visually is the FSIM(c) for its higher correlation with human evaluation than the PSNR [61].

7.3.1 Monochromatic Face Images

Our face denoising experiment uses the Extended Yale-B dataset [62] of 10 different subject, each under 64 different lighting conditions. According to [63], [64], face images of one subject under various illuminations lie approximately on a 9-dimensional subspace, and are therefore suitable for low-rank modeling. We used the pre-cropped 64 images of the first subject and kept them at full resolution. The resulting collection of images constitutes a 3-way tensor of 64 images of size 192×168 . Each mode corresponds respectively to the columns and rows of the images, and to the illumination component. All three are expected to be low-rank due to the spatial correlation within frontal slices and to the correlation between images of the same subject under different illuminations. We present the comparative quantitative performance of the methods tested in Fig. 7, and provide visualizations of the reconstructed first image at the 30 percent noise level in Fig. 8. We report the metrics averaged on the 64

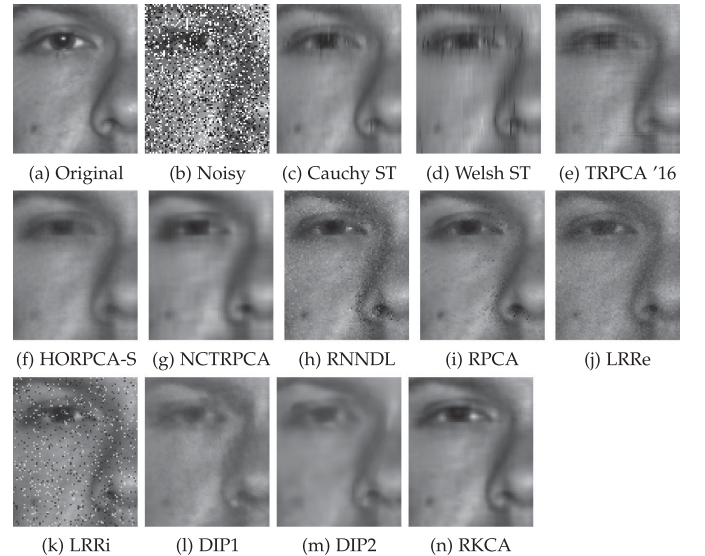






Fig. 8. Results on the Yale benchmark with 30 percent noise. TRPCA '14 removed.

TABLE 2
Three best results on *Yale* at 60 percent Noise.

			
Noisy	RKCA	NC TRPCA	TRPCA '16
PSNR	27.1490	22.8502	22.566
FSIM	0.8913	0.8509	0.8427

images. For DIP, we chose to use a single model for the 64 images, rather than one model per image, to provide a better comparison with the other approaches. In this case, the input dimension was chosen to be 128.

At the 10 percent noise level, nearly every method provided good to excellent recovery of the original images. We therefore omit this noise level (cf. Appendix H, Table 7 and figures, available in the online supplemental material). On the other hand, most methods, with the notable exception of RKCA, NC TRPCA, and TRPCA '16, failed to provide acceptable reconstruction in the gross corruption case. Thus, we present the denoised images at the 30 percent level, and compare the performance of the three best performing methods in Table 2 for the 60 percent noise level.

Clear differences appeared at the 30 percent noise level, as demonstrated both by the quantitative metrics, and by visual inspection of Fig. 8. Overall, performance was markedly lower than at the 10 percent level, and most methods started to lose much of the details. Visual inspection of the results confirms a higher reconstruction quality for RKCA. We invite the reader to look at the texture of the skin, the white of the eye, and at the reflection of the light on the subject's skin and pupil. The latter, in particular, is very close in nature to the white pixel corruption of the salt & pepper noise. Out of all methods, RKCA provided the best reconstruction quality: it is the only algorithm that removed all the noise and for which all the aforementioned details are distinguishable in the reconstruction. Both deep learning approaches provided reconstructions with visually-identifiable features, but under-performed compared to the best robust factorization models.

At the 60 percent noise level, our method scored markedly higher than its competitors on image quality metrics, as seen

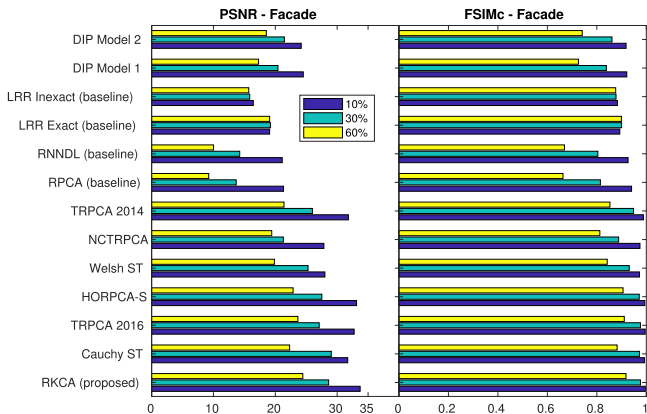


Fig. 9. PSNR and FSIMc of all methods on the Facade benchmark at noise levels 10, 30, and 60 percent.

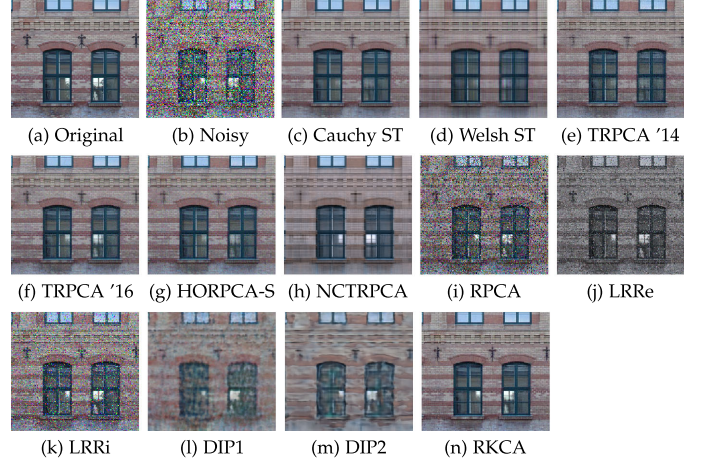


Fig. 10. Results on the Facade benchmark with 30 percent noise.

both in Fig. 7 and in Table 2. Visualizing the reconstructions confirms the difference: the image recovered by RKCA at the 60 percent noise level is comparable to the output of competing algorithms at the 30 percent noise level.

7.3.2 Color Image Denoising

Our benchmark is the *Facade* image [65]: the geometric nature of the building's front wall, and the strong correlation between the RGB bands indicate the data can be modeled by a low-rank 3-way tensor where each frontal slice is a color channel. The rich details and lighting make it interesting to assess fine reconstruction. The input dimension for DIP was set to 32, following the color image denoising examples.


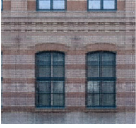
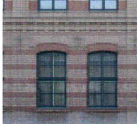
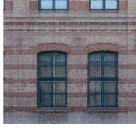
At the 10 percent noise level, RKCA attained the highest PSNR, and the highest FSIMc value. Most methods provided excellent reconstruction, in agreement with the high values of the metrics shown in Fig. 9. As in the previous benchmark, full results are in Appendix H (Table 8 and figures), available in the online supplemental material. At the 30 percent noise level, Cauchy ST exhibited the highest PSNR and RKCA the second highest, while TRPCA '16 and RKCA were tied for first place on the FSIMc metric. Details are provided in Fig. 10. Clear differences in reconstruction quality are visible, and are best seen on the fine details of the picture, such as the black iron ornaments, or the light coming through the window. Our method best preserved the dynamics of the lighting, and the sharpness of the details, and in the end provided the reconstruction visually closest to the original. Competing models tend to oversmooth the image, and to make the light dimmer; indicating substantial losses of high-frequency and dynamic information. RKCA appears to also provide the best color fidelity. Similar to the Yale-B experiment, the deep-learning models are able to handle the gross corruption to a certain extent, but suffer from more distortion than the robust factorizations.

In the gross-corruption case, RKCA was the only method with TRPCA '16 and HORPCA-S to provide a reconstruction with distinguishable details, and did it best (Table 3) while achieving the highest score on both quantitative metrics.

7.4 Tensor Completion

To showcase the tensor completion capabilities of our algorithm, we implemented an LADMM method to solve

TABLE 3
Three Best Results on *Facade* at 60 Percent Noise

			
Noisy	RKCA	TRPCA '16	HORPCA-S
PSNR	24.4491	23.6552	22.8811
FSIMc	0.9179	0.9109	0.9060

problem (24) with $f(\mathcal{L}) = \alpha \|\mathcal{R}\|_1 + \frac{1}{2} (\|\mathbf{A}\|_F^2 + \|\mathbf{B}\|_F^2)$. We provide comparison with one robust tensor completion model (HORPCA-S with missing values [7]) and the matrix Robust PCA with missing values [11].

7.4.1 Yale-B with Noise and Missing Values

Our first experiment extends Section 7.3.1: we investigate the case where apart from corruption, some values are missing. We generated the data by first introducing 30 percent salt & pepper noise, then removing 30 percent of the pixels at random.

As seen in Table 4, RKCA markedly outperformed both other models in terms of FSIM, which translates into a more natural reconstruction. HORPCA-S achieved a similar PSNR but many details are lost, while RPCA removed much of the image's details and left some corruption.

7.4.2 300 Faces in the Wild

Our second experiment is on the completion of unwarped 3D faces with partial self-occlusions taken from the 300 faces in the wild challenge [66], [67], [68].

We present in Table 5 the occluded frames 1, 44, and 76 of a video of the dataset, and the completed frames obtained with RKCA, HORPCA-S, and RPCA. Since no corruption is present in the dataset, the λ parameters were fixed to a high value (1×10^4) such that the algorithms behaved as tensor or matrix completion models. For RKCA, we bounded the rank of the reconstruction to 50 and performed grid-search on 5 values of the α parameter.

It is clear from Table 5 that RKCA provides the most completion. HORPCA-S is able to complete sparse missing values at the center of the image but is lacking on the self-occlusions at the right. Matrix RPCA failed to complete the frames on this benchmark.

TABLE 4
Reconstruction of the First Face of Yale-B with 30 Percent Salt & Pepper Noise and 30 Percent Missing Values







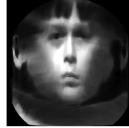
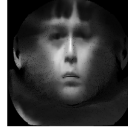



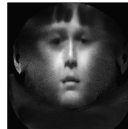
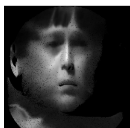

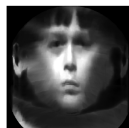
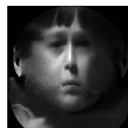
			
Noisy	RKCA	HORPCA-S	RPCA
PSNR	22.7069	22.3827	19.6597
FSIM	0.9332	0.9187	0.8717

TABLE 5
Completion Experiment on the 300W Dataset

			
			
			
Original	RKCA	HORPCA	RPCA

8 CONCLUSION

In this work we presented RKCA, a framework for robust low-rank representation learning on separable dictionaries that expresses the problem as a tensor factorization. We proposed three different regularizers and derived the steps for both ADMM and LADMM solvers. We showed the factorization can be expressed in an equivalent way that gives optimality guarantees when coupled with a regularizer that is positively-homogeneous of degree 3. We then further discussed the low-rank properties of the models, and their practical implementation. We reached the conclusion that our model with a degree 2 regularizer, achieved a good trade-off between experimental effectiveness and parameter tuning difficulty.

Future work can seek to develop a supervised variant of RKCA that could be applied to the same tasks. Another extension that we leave for future work is to consider non-convex regularizers: it is well-known that letting $p \rightarrow 0$ in the element-wise ℓ_p and in the Scatten- p norms recover respectively the ℓ_0 pseudo-norm and the rank functions. We refer to [69] for an example in the matrix RPCA case. Finally, we believe RKCA could be extended with deep learning by replacing the factorization of the low-rank component by a deep neural network, such as an auto-encoder. Effectively, the bilinear factorization we use can be seen as an elementary linear auto-encoder, and replacing it with a more involved non-linear model would be straightforward.

ACKNOWLEDGMENTS

Mehdi Bahri was partially funded by the Department of Computing, Imperial College London. The work of Y. Panagakis has been partially supported by the European Community Horizon 2020 [H2020/2014-2020] under Grant Agreement No. 645094 (SEWA). S. Zafeiriou was partially funded by EPSRC Project EP/N007743/1 (FACER2VM) and also partially funded by a Google Faculty Research Award.

REFERENCES

- [1] K. Pearson, "On lines and planes of closest fit to systems of points in space," *Philosoph. Mag.*, vol. 2, no. 6, pp. 559–572, 1901.

- [2] H. Hotelling, B. A. Olshausen, and D. J. Field, "Analysis of a complex of statistical variables into principal components," *J. Educational Psychology*, vol. 24, no. 6, pp. 417–441, 1933.
- [3] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel, "Backpropagation applied to handwritten zip code recognition," *Neural Comput.*, vol. 1, no. 4, pp. 541–551, Dec. 1989.
- [4] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [5] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, May 2015, Art. no. 436.
- [6] N. Xue, G. Papamakarios, M. Bahri, Y. Panagakis, and S. Zafeiriou, "Robust low-rank tensor modelling using Tucker and CP decomposition," in *Proc. 25th Eur. Signal Process. Conf.*, Sep. 2017, pp. 1185–1189.
- [7] D. Goldfarb and Z. Qin, "Robust low-rank tensor recovery: Models and algorithms," *SIAM J. Matrix Anal. Appl.*, vol. 35, no. 1, pp. 225–253, Jan. 2014.
- [8] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *Proc. IEEE*, vol. 98, no. 6, pp. 1045–1057, Jun. 2010.
- [9] B. A. Olshausen and D. J. Field, "Sparse coding with an overcomplete basis set: A strategy employed by V1?" *Vis. Res.*, vol. 37, no. 23, pp. 3311–3325, Dec. 1997.
- [10] H. Zou, T. Hastie, and R. Tibshirani, "Sparse principal component analysis," *J. Comput. Graph. Statist.*, vol. 15, no. 2, pp. 265–286, 2006.
- [11] E. J. Cands, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, pp. 11:1–11:37, Jun. 2011.
- [12] J. Wright, Y. Ma, J. Mairal, G. Sapiro, T. S. Huang, and S. Yan, "Sparse representation for computer vision and pattern recognition," *Proc. IEEE*, vol. 98, no. 6, pp. 1031–1044, Jun. 2010.
- [13] M. Aharon, M. Elad, and A. Bruckstein, "K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation," *IEEE Trans. Signal Process.*, vol. 54, no. 11, pp. 4311–4322, Nov. 2006.
- [14] J. Mairal, M. Elad, and G. Sapiro, "Sparse representation for color image restoration," *IEEE Trans. Image Process.*, vol. 17, no. 1, pp. 53–69, Jan. 2008.
- [15] G. Liu and P. Li, "Low-rank matrix completion in the presence of high coherence," *IEEE Trans. Signal Process.*, vol. 64, no. 21, pp. 5623–5633, Nov. 2016.
- [16] G. Liu, Q. Liu, and P. Li, "Blessing of dimensionality: Recovering mixture data via dictionary pursuit," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 1, pp. 47–60, Jan. 2017.
- [17] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, "Robust recovery of subspace structures by low-rank representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 171–184, Jan. 2013.
- [18] S. Hawe, M. Seibert, and M. Kleinstueber, "Separable dictionary learning," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognition*, 2013, pp. 438–445.
- [19] S. Hsieh, C. Lu, and S. Pei, "2d sparse dictionary learning via tensor decomposition," in *Signal Information. Process.*, Dec. 2014, pp. 492–496.
- [20] Y. Fang, J. Wu, and B. Huang, "2d sparse signal recovery via 2d orthogonal matching pursuit," *Sci. China Inf. Sci.*, vol. 55, pp. 889–897, 2012.
- [21] F. L. Hitchcock, "The expression of a tensor or a polyadic as a sum of products," *J. Math. Phys.*, vol. 6, no. 1–4, pp. 164–189, Apr. 1927.
- [22] R. Harshman, "Foundations of the PARAFAC procedure: Models and conditions for an" explanatory" multi-modal factor analysis," *UCLA Working Papers Phonetics*, vol. 16, no. 1, 1970, Art. no. 84.
- [23] J. D. Carroll and J.-J. Chang, "Analysis of individual differences in multidimensional scaling via an n-way generalization of Eckart-Young decomposition," *Psychometrika*, vol. 35, no. 3, pp. 283–319, Sep. 1970.
- [24] Z. Shakeri, A. D. Sarwate, and W. U. Bajwa, "Identification of kronecker-structured dictionaries: An asymptotic analysis," in *Proc. IEEE 7th Int. Workshop Computational. Adv. Multi-Sensor Adaptive Process.*, Dec. 2017, pp. 1–5.
- [25] Z. Shakeri, W. U. Bajwa, and A. D. Sarwate, "Minimax lower bounds for Kronecker-structured dictionary learning," in *Proc. IEEE Int. Symp. Information. Theory*, Jul. 2016, pp. 1148–1152.
- [26] Z. Shakeri, W. U. Bajwa, and A. D. Sarwate, "Minimax lower bounds on dictionary learning for tensor data," *IEEE Trans. Inf. Theory*, vol. 64, no. 4, pp. 2706–2726, Apr. 2018.
- [27] Z. Shakeri, W. U. Bajwa, and A. D. Sarwate, "Sample complexity bounds for dictionary learning of tensor data," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, Mar. 2017, pp. 4501–4505.
- [28] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Found. Trends Mach. Learn.*, vol. 3, no. 1, pp. 1–122, Jan. 2011.
- [29] M. Bahri, Y. Panagakis, and S. Zafeiriou, "Robust Kronecker-decomposable component analysis for low-rank modeling," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 3372–3381.
- [30] B. D. Haeffele, E. D. Young, and R. Vidal, "Structured low-rank matrix factorization: Optimality, algorithm and applications to image processing," in *Proc. 31st Int. Conf. Int. Conf. Mach. Learn. - Vol. 32*, pp. II–2007–II–2015.
- [31] B. D. Haeffele and R. Vidal, "Global optimality in tensor factorization, deep learning, and beyond," *CoRR*, vol. abs/1506.07540, 2015. <http://arxiv.org/abs/1506.07540>
- [32] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Deep image prior," in *Proc. IEEE Conf. Pattern Recognit.*, Jun. 2018, pp. 9446–9454.
- [33] T. Kolda and B. Bader, "Tensor decompositions and applications," *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, Aug. 2009.
- [34] Z. Zhang, G. Ely, S. Aeron, N. Hao, and M. Kilmer, "Novel methods for multilinear data completion and de-noising based on tensor-SVD," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 3842–3849.
- [35] C. Lu, J. Feng, Y. Chen, W. Liu, Z. Lin, and S. Yan, "Tensor robust principal component analysis: Exact recovery of corrupted low-rank tensors via convex optimization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2016, pp. 5249–5257.
- [36] F. Shang, Y. Liu, J. Cheng, and H. Cheng, "Robust principal component analysis with missing data," in *Proc. 23rd ACM Int. Conf. Conf. Inf. Knowl. Manage.*, 2014, pp. 1149–1158.
- [37] A. J. Laub, *Matrix Analysis For Scientists And Engineers*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2004.
- [38] B. Recht, M. Fazel, and P. Parrilo, "Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization," *SIAM Rev.*, vol. 52, no. 3, pp. 471–501, Jan. 2010.
- [39] G. Golub, S. Nash, and C. Van Loan, "A Hessenberg-Schur method for the problem $AX + XB = C$," *IEEE Trans. Automat. Control*, vol. 24, no. 6, pp. 909–913, Dec. 1979.
- [40] X. Peng, C. Lu, Z. Yi, and H. Tang, "Connections between nuclear-norm and frobenius-norm-based representations," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 1, pp. 218–224, Jan. 2018.
- [41] H. Zhang, Z. Yi, and X. Peng, "FIRR: Fast low-rank representation using Frobenius-norm," *Electron. Lett.*, vol. 50, no. 13, pp. 936–938, Jun. 2014.
- [42] C. Grussler, A. Rantzer, and P. Giselsson, "Low-rank optimization with convex constraints," *IEEE Trans. Autom. Control*, vol. 63, no. 11, pp. 4000–4007, Nov. 2018.
- [43] Z. Lin, R. Liu, and Z. Su, "Linearized alternating direction method with adaptive penalty for low-rank representation," in *Proc. 24th Int. Conf. Neural Inf. Process. Syst.*, 2011, pp. 612–620.
- [44] Z. Xu, M. Figueiredo, and T. Goldstein, "Adaptive ADMM with spectral penalty parameter selection," in *Proc. 20th Int. Conf. Artif. Intell. Statist.*, Apr. 2017, pp. 718–727.
- [45] Y. Xu, M. Liu, Q. Lin, and T. Yang, "ADMM without a fixed penalty parameter: Faster convergence with new adaptive penalization," in *Proc. Neural Inf. Process. Syst.*, 2017, pp. 1267–1277.
- [46] M. Hong, Z. Luo, and M. Razaviyayn, "Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems," *SIAM J. Optimization*, vol. 26, no. 1, pp. 337–364, Jan. 2016.
- [47] Y. Wang, W. Yin, and J. Zeng, "Global convergence of ADMM in nonconvex nonsmooth optimization," *J. Sci. Comput.*, vol. 78, no. 1, pp. 29–63, Jan. 2019. <https://doi.org/10.1007/s10915-018-0757-z>
- [48] T. F. Chan, "An improved algorithm for computing the singular value decomposition," *ACM Trans. Math. Softw.*, vol. 8, no. 1, pp. 72–83, Mar. 1982.
- [49] J. Dean and S. Ghemawat, "MapReduce: Simplified data processing on large clusters," in *Operating Syst. Implementation*, 2004, pp. 137–150.
- [50] X. Peng, J. Lu, Z. Yi, and R. Yan, "Automatic subspace learning via principal coefficients embedding," *IEEE Trans. Cybern.*, vol. 47, no. 11, pp. 3583–3596, Nov. 2017.
- [51] Z. Lin, R. Liu, and S. Zhixun, "Linearized Alternating Direction Method with Adaptive Penalty for Low-Rank Representation," *Advances in Neural Information Processing Systems 24*, J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira and K. Q. Weinberger, Eds., Curran Associates, Inc., 2011, pp. 612–620. <http://papers.nips.cc/paper/4434-linearized-alternating-direction-method-with-adaptive-penalty-for-low-rank-representation.pdf>

- [52] Q. Pan, D. Kong, C. Ding, and B. Luo, "Robust non-negative dictionary learning," in *Proc. 28th AAAI Conf. Artif. Intell.*, 2014, pp. 2027–2033.
- [53] Y. Yang, Y. Feng, and J. A. K. Suykens, "Robust low-rank tensor recovery with regularized redescending M-Estimator," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 27, no. 9, pp. 1933–1946, Sep. 2016.
- [54] A. Anandkumar, P. Jain, Y. Shi, and U. N. Niranjan, "Tensor vs. matrix methods: Robust tensor decomposition under block sparse perturbations," in *Proc. 19th Int. Conf. Artif. Intell. Statist.*, 2016, pp. 268–276.
- [55] P. Netrapalli, N. U. N. S. Sanghavi, A. Anandkumar, and P. Jain, "Non-convex robust PCA," in *Proc. 27th Int. Conf. Neural Inf. Process. Syst.* - Vol. 1, 2014, pp. 1107–1115.
- [56] N. Goyette, P. Jodoin, F. Porikli, J. Konrad, and P. Ishwar, "Changetection.net: A new change detection benchmark dataset," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2012, pp. 1–8.
- [57] L. Li, W. Huang, I. Y.-H. Gu, and Q. Tian, "Statistical modeling of complex backgrounds for foreground object detection," *IEEE Trans. Image Process.*, vol. 13, no. 11, pp. 1459–1472, Nov. 2004.
- [58] Q. Zhao, G. Zhou, L. Zhang, A. Cichocki, and S. Amari, "Bayesian robust tensor factorization for incomplete multiway data," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 27, no. 4, pp. 736–748, Apr. 2016.
- [59] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, 2006.
- [60] L. Zhang, L. Zhang, X. Mou, and D. Zhang, "FSIM: A feature similarity index for image quality assessment," *IEEE Trans. Image Process.*, vol. 20, no. 8, pp. 2378–2386, Aug. 2011.
- [61] P. Mohammadi, A. Ebrahimi-Moghadam, and S. Shirani, "Subjective and objective quality assessment of image: A survey," *Majlesi J. Electr. Eng.*, vol. 9, no. 1, Dec. 2014, pp. 55–83. Retrieved from <http://mjee.iaumajlesi.ac.ir/index/index.php/ee/article/view/1376>
- [62] A. S. Georgiades, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 643–660, Jun. 2001.
- [63] R. Ramamoorthi and P. Hanrahan, "An efficient representation for irradiance environment maps," in *Proc. 28th Annu. Conf. Comput. Graph. Interactive Tech.*, 2001, pp. 497–500.
- [64] R. Basri and D. W. Jacobs, "Lambertian reflectance and linear subspaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 2, pp. 218–233, Feb. 2003.
- [65] X. Chen, Z. Han, Y. Wang, Q. Zhao, D. Meng, and Y. Tang, "Robust tensor factorization with unknown noise," in *Proc. IEEE Conf. Computer. Vis. Pattern Recognit.*, Jun. 2016, pp. 5213–5221.
- [66] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "300 faces in-the-wild challenge: The first facial landmark localization challenge," in *Proc. IEEE Int. Conf. Comput. Vis. Workshops*, Dec. 2013, pp. 397–403.
- [67] C. Sagonas, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "A semi-automatic methodology for facial landmark annotation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops*, Jun. 2013, pp. 896–903.
- [68] C. Sagonas, E. Antonakos, G. Tzimiropoulos, S. Zafeiriou, and M. Pantic, "300 faces in-the-wild challenge: Database and results," *Image Vis. Comput.*, vol. 47, pp. 3–18, 2016.
- [69] G. Papamakarios, Y. Panagakis, and S. Zafeiriou, "Generalised scalable robust principal component analysis," in *Proc. British Mach. Vis. Conf.*, M.Valstar, A.French and T.Pridmore, Eds., BMVA Press, 2014. doi: <http://dx.doi.org/10.5244/C.28.116>



Mehdi Bahri received the diplôme d'Ingénieur in applied mathematics with honours from Grenoble Institute of Technology - Ensimag, and the MSc degree in advanced computing with distinction from Imperial College London, both in 2016. He is working toward the PhD degree with Stefanos Zafeiriou in the Department of Computing, Imperial College London. He published his research in ICCV and spent one year in the industry. His research interests include representation learning, geometric deep learning, and Bayesian learning. He is a student member of the IEEE.



Yannis Panagakis received the BSc degree in informatics and telecommunication from the University of Athens, Greece, and the MSc and PhD degrees from the Department of Informatics, Aristotle University of Thessaloniki. He is a senior lecturer (associate professor equivalent) in computer science with Middlesex University London and a research fellow with the Department of Computing, Imperial College London. His research interests include machine learning and its interface with signal processing, high-dimensional statistics, and computational optimization. Specifically, he is working on models and algorithms for robust and efficient learning from high-dimensional data and signals representing audio, visual, affective, and social information. He has been awarded the prestigious Marie-Curie Fellowship, among various scholarships and awards for his studies and research. He currently serves as the managing editor of the *Image and Vision Computing Journal*. He co-organized the BMVC 2017 conference and several workshops and special sessions in top venues such as ICCV.



Stefanos Zafeiriou (M'09) is currently a reader in machine learning and computer vision with the Department of Computing, Imperial College London, London, United Kingdom, and a distinguished research fellow with the University of Oulu under the Finnish Distinguishing Professor Programme. He was a recipient of the Prestigious Junior Research Fellowships from Imperial College London in 2011 to start his own independent research group. He was the recipient of the Presidents Medal for Excellence in Research Supervision for 2016. He is recipient of many best paper awards including the best student paper award in FG'2018. In 2018, he received an EPSRC Fellowship. He currently serves as an associate editor of the *IEEE Transactions on Affective Computing* and *Computer Vision* and the *Image Understanding* journal. In the past he held editorship positions in the *IEEE Transactions on Cybernetics*, the *Image and Vision Computing Journal*. He has been a guest editor of over six journal special issues and co-organised over 13 workshops/special sessions on specialised computer vision topics in top venues, such as CVPR/FG/ICCV/ECCV (including three very successfully challenges run in ICCV'13, ICCV'15 and CVPR'17 on facial landmark localisation/tracking). He has co-authored more than 60 journal papers mainly on novel statistical machine learning methodologies applied to computer vision problems, such as 2-D/3-D face analysis, deformable object fitting and tracking, shape from shading, and human behaviour analysis, published in the most prestigious journals in his field of research, such as the *IEEE Transactions on Pattern Analysis and Machine Intelligence*, the *International Journal of Computer Vision*, the *IEEE Transactions on Image Processing*, the *IEEE Transactions on Neural Networks and Learning Systems*, the *IEEE Transactions on Visualization and Computer Graphics*, and the *IEEE Transactions on Information Forensics and Security*, and many papers in top conferences, such as CVPR, ICCV, ECCV, ICML. His students are frequent recipients of very prestigious and highly competitive fellowships, such as the Google Fellowship x2, the Intel Fellowship, and the Qualcomm Fellowship x3. He has more than 7200 citations to his work, h-index 44. He was the general chair of BMVC 2017. He is a member of the IEEE.

► For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.