Robust Low-Rank Tensor Recovery With Regularized Redescending M-Estimator

Yuning Yang, Yunlong Feng, and Johan A. K. Suykens, *Fellow, IEEE*

***Abstract*— This paper addresses the robust low-rank tensor recovery problems. Tensor recovery aims at reconstructing a low-rank tensor from some linear measurements, which finds applications in image processing, pattern recognition, multitask learning, and so on. In real-world applications, data might be contaminated by sparse gross errors. However, the existing approaches may not be very robust to outliers. To resolve this problem, this paper proposes approaches based on the regularized redescending M-estimators, which have been introduced in robust statistics. The robustness of the proposed approaches is achieved by the regularized redescending M-estimators. However, the nonconvexity also leads to a computational difficulty. To handle this problem, we develop algorithms based on proximal and linearized block coordinate descent methods. By explicitly deriving the Lipschitz constant of the gradient of the data-fitting risk, the descent property of the algorithms is present. Moreover, we verify that the objective functions of the proposed approaches satisfy the Kurdyka–Łojasiewicz property, which establishes the global convergence of the algorithms. The numerical experiments on synthetic data as well as real data verify that our approaches are robust in the presence of outliers and still effective in the absence of outliers.**

***Index Terms*— Block coordinate descent, global convergence, nonconvexity, redescending M-estimator, robust tensor recovery.**

I. Introduction

TENSORS, appearing as the higher order generalization of vectors and matrices, make it possible to represent data that have intrinsically many dimensions, and give a better understanding of the relationship behind the information from a higher order perspective. Similar to matrices, some higher order data may naturally have sparse structure in some sense, e.g., the low-rank property, or include dominant information in a few factors. To explore the underlying structure of higher order tensors, two decomposition strategies, namely: 1) the Tucker decomposition and 2) the CP decomposition [1]–[4] are commonly used.

Although the above two decomposition strategies can decompose or approximate the underlying structure of tensors, the fact that they are built under the framework of unsupervised learning limits their use when some information is known *a priori*. On the other hand, the strategies may be sensitive to outliers. Recently, within the supervised learning and convex optimization frameworks, Liu *et al.* [5] and Gandy *et al.* [6] independently proposed approaches for recovering a low-rank tensor from some linear measurements, among which recovering a low-rank tensor from incomplete observations is of particular interest. These new approaches take some prior knowledge into consideration, and moreover, they introduce the tensor nuclear norm to model the low-rank property of tensors. Applications, including image processing [7], pattern recognition [8], multitask learning [9], spectral data [10], and intelligent transportation systems [11], demonstrate the effectiveness of the new approaches. Apart from these applications, other research are carried out on understanding the theoretical behaviors and improving the computational efficiency of the new approaches and their variants (see [12]–[16]).

In the presence of noise, the low-rank tensor recovery approaches incorporate the learning framework of loss + regularization, where the least squares loss is commonly adopted (see [6], [13]). It is known that the least squares loss possesses the interpretation of the maximum likelihood estimation when the distribution of the noise is Gaussian. However, in real-world applications, such as image processing [17], the distribution of the noise is usually unknown, and data might be grossly corrupted by outliers, whereas in this case, the least squares loss-based approaches behave poorly due to their nonrobustness.

To address tensor recovery problems with outliers, robust approaches were proposed in [18]–[20]. Specifically, [18] and [19] focused on the tensor principal component analysis (PCA) problems, whereas [20] took tensor completion as well as tensor PCA into account, the robustness of which benefits from a robust loss function—the least absolute deviations (LAD) loss $|\cdot|$. The LAD loss shares the ability of being relatively insensitive to large deviations in outlying observations, and its convexity also leads to the global solution to the above approaches. However, the LAD loss still penalizes outliers linearly, which may not be very robust [21]. Inoue *et al.* [22] considered robust multilinear PCA using the Welsch loss, which is in fact a redescending M-estimator. As proposed in

Manuscript received May 19, 2014; revised January 15, 2015, May 11, 2015, and July 23, 2015; accepted July 30, 2015. Date of publication August 19, 2015; date of current version August 15, 2016. This work was supported in part by the European Research Council under the European Union's Seventh Framework Programme under Grant FP7/2007-2013 and Grant ERC AdG A-DATADRIVE-B (290923), in part by the Research Council KUL under Grant GOA/10/09 MaNet, Grant CoE PFV/10/002 (OPTEC), and Grant BIL12/11T, in part by PhD/Postdoc grants, in part by Flemish Government: FWO: PhD/Postdoc grants under Project G.0377.12 (structured systems) and Project G.088114N (tensor-based data similarity), in part by IWT: PhD/Postdoc grants under Project SBO POM (100031), in part by iMinds Medical Information Technologies under Grant SBO 2014, and in part by Belgian Federal Science Policy Office under Grant IUAP P7/19 (dynamical systems, control and optimization, 2012–2017).

The authors are with the Department of Electrical Engineering, Stadius Centre for Dynamical Systems, Signal Processing and Data Analytics, Katholieke Universiteit Leuven, Leuven B-3001, Belgium (e-mail: yuning.yang@esat.kuleuven.be; yunlong.feng@esat.kuleuven.be; Johan.Suykens@esat.kuleuven.be).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/TNNLS.2015.2465178

2162-237X © 2015 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

robust statistics [21], redescending M-estimators have been widely used in robust regression problems. The redescending property makes the redescending M-estimators be particularly resistant to extreme outliers and may outperform bounded M-estimators [23], such as those induced by the LAD loss.

Inspired by the success of applications of redescending M-estimators, in this paper, the Welsch loss and the Cauchy loss-based redescending M-estimators are employed to seek robustness when recovering a low-rank tensor from observations contaminated by non-Gaussian noise or outliers. Our considerations of using these two losses are also motivated by the theoretical investigations presented in [24]–[26] and the empirical successes reported in [22] and [27]–[29]. To explore the underlying low-rank tensor data, the mixture tensor model [12], [13], which is the sum of some factor tensors and each of which is low rank in only a specific mode, is adopted. Then, the nuclear norm regularization and the rank constrained strategies are incorporated to ensure the low-rank property of each factor.

By exploiting the separable structure of the proposed models, proximal and linearized block coordinate descent methods are developed. To show the descent property of the algorithms, the Lipschitz constant of the gradient of the data-fitting risk is derived. Moreover, we verify the analytic property of the data-fitting risk and the semialgebraic property of the regularization schemes, which, within the framework of [30], show the global convergence of the developed algorithms.

The remainder of this paper is organized as follows. In Section II, we introduce preliminaries on basic tensor algebra and tensor recovery problems. We propose our robust tensor recovery approaches based on the regularized redescending M-estimators in Section III and give some discussions. We then provide a proximal and linearized block coordinate descent method in Section IV, and verify their convergence properties in Section V. The performance of our methods is tested in Section VI on synthetic data as well as real data. The conclusion is drawn in Section VII.

II. Preliminaries

A. Basic Tensor Algebra

A tensor is a multiway array. An Nth-order tensor is denoted as $\mathcal{X} \in \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_N}$. We use $\mathbb{T} := \mathbb{R}^{n_1 \times n_2 \times \cdots \times n_N}$ to denote an Nth-order tensor space throughout this paper. Below we list some concepts of tensors. The readers can refer to [31] for more details.

1) Tensor-Matrix Unfolding: A fiber of a tensor $\mathcal{X}$ is a column vector obtained by fixing every index of $\mathcal{X}$ but one. The mode-d unfolding, or matricization, denoted by $\mathbf{X}_{(d)}$, is a matrix of size $n_d \times \prod_{i \neq d}^{N} n_i$, whose columns are the mode-d fibers of $\mathcal{X}$ in the lexicographical order.

2) Inner Product and Tensor-Matrix Multiplication: For $\mathcal{A}, \mathcal{B} \in \mathbb{T}$, their inner product is given by

$$\langle \mathcal{A}, \mathcal{B} \rangle = \sum_{i_1=1}^{n_1} \cdots \sum_{i_N=1}^{n_N} \mathcal{A}_{i_1 \ldots i_N} \mathcal{B}_{i_1 \ldots i_N}.$$

The Frobenius norm of $\mathcal{A}$ is defined by $\|\mathcal{A}\|_F = \langle \mathcal{A}, \mathcal{A} \rangle^{1/2}$. The mode-$d$ tensor-matrix multiplication of a tensor $\mathcal{X} \in \mathbb{T}$ with a matrix $\mathbf{U} \in \mathbb{R}^{J_d \times n_d}$, denoted by $\mathcal{Y} = \mathcal{X} \times_d \mathbf{U}$, is of size $n_1 \times \cdots \times n_{d-1} \times J_d \times n_{d+1} \times \cdots \times n_N$, and can be expressed in terms of unfolding as $\mathbf{Y}_{(d)} = \mathbf{U}\mathbf{X}_{(d)}$.

3) Tensor Rank and Decompositions: There are mainly two types of tensor rank, namely: 1) the CP-rank and 2) the mode-d rank. The CP-rank is defined as the minimum positive integer R such that for a tensor $\mathcal{X} \in \mathbb{T}$, it can be factorized as a sum of R rank-one tensors. This factorization is called the CP decomposition. Finding such a decomposition exactly is NP-hard for higher order tensors [31]. The mode-d rank of a tensor $\mathcal{X} \in \mathbb{T}$, also known as the Tucker rank, is defined as the rank of the mode-d unfolding matrix $\mathbf{X}_{(d)}$. $\mathcal{X}$ is said to be rank-$(R_1, \ldots, R_N)$ if the mode-d rank of $\mathcal{X}$ is $(R_1, \ldots, R_N)$. Any rank-$(R_1, \ldots, R_N)$ tensor $\mathcal{X}$ can be decomposed as $\mathcal{X} = \mathcal{G} \times_1 \mathbf{U}^1 \times_2 \cdots \times_N \mathbf{U}^N$, where $\mathcal{G} \in \mathbb{R}^{R_1 \times \cdots \times R_N}$ is called the core tensor and $\mathbf{U}^i \in \mathbb{R}^{n_i \times R_i}$ are factor matrices. This decomposition is called the Tucker decomposition.

4) Tuple of Tensors: Given a set of tensors $\mathcal{X}_1, \ldots, \mathcal{X}_N \in \mathbb{T}$, we denote the tuple of tensors $\boldsymbol{\mathcal{X}}$ as

$$\boldsymbol{\mathcal{X}} := \{\mathcal{X}_1, \mathcal{X}_2, \ldots, \mathcal{X}_N\}.$$

Similar to the notation for vector spaces, we denote the space of tuples of tensors as $\mathbb{T}^N$. For ease of notation, the mode-l unfolding of the kth tensor is denoted as $\mathbf{X}_{k,(l)}$. For convenience, we also define the summation operator $\sum : \mathbb{T}^N \to \mathbb{T}$ such that $\sum(\boldsymbol{\mathcal{X}}) = \sum_{i=1}^{N} \mathcal{X}_i$.

B. Tensor Recovery Problems

The goal of tensor recovery is to find a tensor $\mathcal{X} \in \mathbb{T}$ satisfying $\boldsymbol{\mathcal{A}}(\mathcal{X}) = \mathbf{b}$, where $\mathbf{b} \in \mathbb{R}^p$ is a vector, and $\boldsymbol{\mathcal{A}} : \mathbb{T} \to \mathbb{R}^p$ with $p \leq \prod_{i=1}^{N} n_i$ is a linear map defined as

$$\boldsymbol{\mathcal{A}}(\cdot) = [\langle \mathcal{A}_1, \cdot \rangle, \langle \mathcal{A}_2, \cdot \rangle, \ldots, \langle \mathcal{A}_p, \cdot \rangle]^T$$

with $\mathcal{A}_i \in \mathbb{T}, 1 \leq i \leq p$. Along with the low-rank requirement on $\mathcal{X}$, the low-rank tensor recovery problem reads as follows:

$$\text{finding a low-rank tensor } \mathcal{X} \quad \text{s.t. } \boldsymbol{\mathcal{A}}(\mathcal{X}) = \mathbf{b}.$$

Since the mode-d rank is easier to compute than the CP-rank, the problem boils down to [6]

$$\min_{\mathcal{X} \in \mathbb{T}} \sum_{i=1}^{N} \text{rank}(\mathbf{X}_{(i)}) \quad \text{s.t. } \boldsymbol{\mathcal{A}}(\mathcal{X}) = \mathbf{b}. \tag{1}$$

A special and important case of (1) is the tensor completion problem, which aims at recovering a low-rank tensor from partial observations

$$\min_{\mathcal{X} \in \mathbb{T}} \sum_{i=1}^{N} \text{rank}(\mathbf{X}_{(i)}) \quad \text{s.t. } \mathcal{X}_{i_1 \ldots i_N} = \mathcal{B}_{i_1 \ldots i_N}, \quad (i_1 \ldots i_N) \in \Omega$$

where $\mathcal{B} \in \mathbb{T}$ represents the observed tensor and Ω denotes the set of multi-indices that correspond to the observed entries. To get a tractable problem, the objective functions in the above two problems are replaced by the sum of nuclear norms, which is also called the tensor nuclear norm (see [5], [6], [9]).

So far, we have reviewed the noiseless tensor recovery problems. Suppose now the observation is given by

$\mathbf{b} = \mathcal{A}(\mathcal{X}) + \boldsymbol{\epsilon}$, where $\boldsymbol{\epsilon} \in \mathbb{R}^p$ represents noise or outliers. A common approach to handle the above problem takes the form $\min_{\mathcal{X}\in\mathbb{T}} f(\mathcal{X}) + g(\mathcal{X})$, where f refers to the empirical risk while g is the regularization term, such as the tensor nuclear norm regularization. In the presence of Gaussian noise, the least squares loss-based approaches have been frequently adopted (see [6], [13], [32]). With the presence of gross errors or outliers, robust techniques should be adopted. Based on the LAD loss, the approaches for the tensor PCA and tensor completion problems have been proposed in [18]–[20]. In particular, Goldfarb and Qin [20] introduced the LAD loss into different tensor models and provided a variety of approaches. An important advantage of the LAD loss-based approaches is the convexity. However, as we have pointed out, the LAD loss-based approaches may not be very robust, as they penalize outliers linearly.

III. Tensor Recovery With Regularized Redescending M-Estimator and Discussion

A. Proposed Approaches

Similar to [12] and [13], we assume that the tensor to be recovered can be approximated by a sum of N factor tensors of the same size, each being low rank in only one mode. Specifically, we let $\mathcal{X} := \sum(\boldsymbol{\mathcal{X}}) = \sum_{j=1}^{N} \mathcal{X}_j$, with $\mathbf{X}_{j,(j)}$ being a low rank matrix, $1 \le j \le N$, where $\mathbf{X}_{j,(j)}$ represents the mode-j unfolding matrix of the jth factor tensor. This is called the mixture model [12], [13]. Based on this representation, our data-fitting risk takes the following form:

$$J_\sigma(\textstyle\sum(\boldsymbol{\mathcal{X}})) := \displaystyle\sum_{i=1}^{p} \rho_\sigma(\langle \mathcal{A}_i, \textstyle\sum(\boldsymbol{\mathcal{X}})\rangle - \mathbf{b}_i)$$

where $\sigma > 0$ is a scale parameter and ρ_σ is a robust loss whose influence function has the redescending property. Specifically, let $\psi_\sigma(t) = \rho'_\sigma(t)$ be the empirical influence function. According to [21], ψ_σ of a redescending M-estimator satisfies the property that $\lim_{|t|\to+\infty} \psi_\sigma(t) = 0$. In this paper, we employ the following two losses.

1) *Cauchy Loss:* $\rho_\sigma(t) = 0.5\sigma^2 \log(1 + t^2/\sigma^2)$.
2) *Welsch Loss:* $\rho_\sigma(t) = 0.5\sigma^2(1 - \exp(-t^2/\sigma^2))$.

For tensor completion problems, J_σ can be simplified as

$$J_\sigma(\textstyle\sum(\boldsymbol{\mathcal{X}})) = \displaystyle\sum_{(i_1\dots i_N)\in\Omega} \rho_\sigma\left(\textstyle\sum(\boldsymbol{\mathcal{X}})_{i_1\dots i_N} - \mathcal{B}_{i_1\dots i_N}\right).$$

Concerning the regularization terms and considering the low-rank property of $\mathbf{X}_{j,(j)}$, motivated by [12] and [13], we can use the nuclear norm to penalize these unfolding matrices. Thus, the approach based on the regularized redescending M-estimator is formulated as

$$\min_{\boldsymbol{\mathcal{X}}\in\mathbb{T}^N} \; J_\sigma\left(\sum_{j=1}^{N} \mathcal{X}_j\right) + \lambda \sum_{j=1}^{N} \|\mathbf{X}_{j,(j)}\|_* \tag{2}$$

where $\lambda > 0$ is a regularization parameter.

We also propose the following rank constrained redescending M-estimator-based approach, when the rank information of the tensor to be recovered is known in advance:

$$\min_{\boldsymbol{\mathcal{X}}\in\mathbb{T}^N} \; J_\sigma\left(\sum_{j=1}^{N} \mathcal{X}_j\right) \quad \text{s.t. } \operatorname{rank}(\mathbf{X}_{j,(j)}) \le R_j, \; 1 \le j \le N. \tag{3}$$

It is known that (3) can be reformulated as an unconstrained problem by introducing the indicator functions. Let C be a nonempty and closed set. The indicator function of C, denoted as δ_C, is defined by

$$\delta_C(x) = \begin{cases} 0, & \text{if } x \in C \\ +\infty, & \text{otherwise.} \end{cases}$$

For each $j = 1, \dots, N$, denote

$$M_{R_j} = \{\mathbf{X} \in \mathbb{R}^{n_j \times \prod_{k\neq j}^{N} n_k} \mid \operatorname{rank}(\mathbf{X}) \le R_j\}.$$

Then, M_{R_j} is a nonempty and closed set for each j. Using the indicator functions, (3) can be reformulated as

$$\min_{\boldsymbol{\mathcal{X}}\in\mathbb{T}^N} \; J_\sigma\left(\sum_{j=1}^{N} \mathcal{X}_j\right) + \sum_{j=1}^{N} \delta_{M_{R_j}}(\mathbf{X}_{j,(j)}). \tag{4}$$

For convenience, in the rest of this paper, if necessary, we unify the regularizations $\lambda\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$ by the notation $g(\cdot)$. We also denote

$$\Phi(\boldsymbol{\mathcal{X}}) = \Phi(\mathcal{X}_1, \dots, \mathcal{X}_N) := J_\sigma\left(\sum_{j=1}^{N} \mathcal{X}_j\right) + \sum_{j=1}^{N} g(\mathbf{X}_{j,(j)}) \tag{5}$$

and then the models (2) and (4) can be unified as

$$\min_{\boldsymbol{\mathcal{X}}\in\mathbb{T}^N} \Phi(\boldsymbol{\mathcal{X}}). \tag{6}$$

Remark 1: The matrix Schatten-p norm regularization has been used for matrix recovery [33]. The advantage of the Schatten-p norm is that can approximate the rank function when $p \to 0$. Thus, it will be also interesting to replace the nuclear norm in (2) by the Schatten-p norm.

B. Tuning the Regularization Parameters

In the proposed approaches, the regularization parameters are important in controlling the low-rank property of the resulting tensors. We first consider the λ value of (2). In [20] and [34], the value $(n_{\max})^{1/2}$ is suggested for the regularization parameter of the LAD loss-based matrix and tensor completion approaches, where $n_{\max}$ denotes the largest dimension of a tensor. Motivated by this, in this paper, we empirically find that $\lambda = \alpha n_{\min}/(n_{\max})^{1/2}$ is a proper choice, where $\alpha > 0$ is to be tuned, usually in the interval $(0, 0.5)$, and $n_{\min}$ denotes the smallest dimension of a tensor.

The ranks $R_1, \dots, R_N$ in (3) or (4) can be regarded as regularization parameters. If we know the ranks of the tensor to be recovered, then we can set the values of $R_1, \dots, R_N$ accordingly. Otherwise, the values of the ranks have also to be tuned. A possible and simple heuristic is to apply some algorithms to solve (2) at first with finite steps, and then use the generated rank information in (3), possibly rescaled

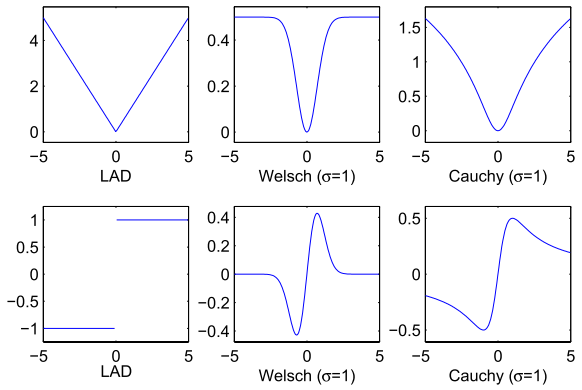


Fig. 1. Plots of different losses (top) and their influence functions (bottom).

by a factor smaller than 1. In the PCA setting, i.e., $\mathcal{A}$ is an identity, a strategy similar to the Q-based method for multilinear PCA [35] can be applied. In our setting, we define $Q^{(d)} = \sum_{i_d=1}^{R_d} \lambda_{i_d}^{(d)} / \sum_{i_d=1}^{n_d} \lambda_{i_d}^{(d)}$, where $\lambda_{i_d}^{(d)}$ is the i-th singular value of the mode-d unfolding matrix $B_{(d)}$, with $\lambda_{1_d}^{(d)} \geq \lambda_{2_d}^{(d)} \geq \cdots \geq \lambda_{n_d}^{(d)}$, $1 \leq i_d \leq n_d$, $1 \leq d \leq N$. One then chooses some value $\epsilon \in (0, 1)$ and selects R_d such that $Q^{(d)} \approx \epsilon$ for $d = 1, \ldots, N$. Thus, by choosing a suitable ϵ, one ignores the less important eigenvectors of the unfoldings of $\mathcal{B}$. Other heuristics are also available in [36].

C. Discussions on the Employed Loss Functions

This part presents some discussions on the employed loss functions. To make the discussions convenient, we first distinguish the two employed losses by $\rho_\sigma^{\text{Welsh}}$ and $\rho_\sigma^{\text{Cauchy}}$, respectively. We have the following observations.

1) From their corresponding influence functions

$$\psi_\sigma^{\text{Welsh}}(t) = \exp(-t^2/\sigma^2)t$$
$$\psi_\sigma^{\text{Cauchy}}(t) = t/(1 + t^2/\sigma^2).$$

One can see that they both satisfy $\lim_{|t|\to+\infty} \psi_\sigma(t) = 0$. As a result, associated with the robust loss functions $\rho_\sigma^{\text{Welsh}}(t)$ and $\rho_\sigma^{\text{Cauchy}}(t)$, the proposed approaches (2) and (4) are essentially regularized redescending M-estimators. In fact, the value $\psi(t)/t$ can be regarded as a weight of t. One observes that for redescending M-estimators, as t increases, $\psi(t)/t$ decreases sharply, which gives a small weight to the value t. On the other hand, the influence function of the LAD loss is only bounded instead of redescending, and the LAD loss penalizes the large deviation linearly. This leads to the result that (2) and (4) are more robust than the LAD-based approaches. To give an intuitive impression, we plot the three losses and their influence functions in Fig. 1, from which we can easily see the redescending property of ψ^{Welsh} and ψ^{Cauchy}.

2) Another nice property of the influence functions of Welsh loss and Cauchy loss is the Lipschitz continuity, i.e., for any $t_1, t_2 \in \mathbb{R}$, $|\psi_\sigma(t_1) - \psi_\sigma(t_2)| \leq |t_1 - t_2|$. This property is very important and serves as a basic for the convergence of the algorithms presented later.

3) As $t \to 0$, $\rho_\sigma^{\text{Welsh}}$ and $\rho_\sigma^{\text{Cauchy}}$ approximate the least squares loss, which can be seen from their Taylor series. Given a fixed $\sigma > 0$, both of the two losses possess the form $\rho_\sigma(t) = t^2/2 + o((t/\sigma)^2)$. Therefore, $\rho_\sigma(t) \approx t^2/2$ provided $t/\sigma \to 0$. This also reminds us that a large σ can lead to the closeness between ρ_σ and the least squares loss. Such a property gives more flexibility to the two employed losses than LAD. We also mention that a comparative study on the parameter σ of the Welsch loss can be found in [26].

4) Although the employed losses enjoy nice robustness property, their nonconvexity seems to be a barrier in practical use. Nevertheless, from experiments we find that approaches based on the employed losses can yield satisfying results, as will be shown later.

5) Other redescending type losses, such as the Tukey's biweight loss, can also be considered.

IV. Proximal and Linearized Block Coordinate Descent Method

This part concerns with algorithms for (2) and (4). Classical robust M-estimators can be solved by the iteratively reweighted least squares (IRLS) algorithms. However, due to the existence of regularization terms, IRLS cannot be directly applied to solve (2) and (4), and the special structure of Φ requires us to seek other proper methods. In the following, a block coordinate descent type method is developed. First, we discuss some properties of $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$.

A. Properties of $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$

First, we show that $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$ are *proper* and *lower semicontinuous*, which are essential in deriving the Moreau envelopes and are important in the convergence analysis of the developed algorithms. The definitions are given as follows.

Definition 1 (See [37]): A function f is called *lower semicontinuous* at $\mathbf{x}_0 \in \mathbb{R}^n$, if for every $\epsilon > 0$, there is a neighborhood $\mathbb{U}(\mathbf{x}_0, \epsilon)$ of $\mathbf{x}_0$ such that for all $\mathbf{x} \in \mathbb{U}(\mathbf{x}_0, \epsilon)$, there holds $f(\mathbf{x}) \geq f(\mathbf{x}_0) - \epsilon$. Equivalently, $\liminf_{\mathbf{x}\to\mathbf{x}_0} f(\mathbf{x}) \geq f(\mathbf{x}_0)$. f is called *proper* if $f(\mathbf{x}) < \infty$ for at least one $\mathbf{x} \in \mathbb{R}^n$, and $f(\mathbf{x}) > -\infty$ for all $\mathbf{x} \in \mathbb{R}^n$.

Proposition 1: Both $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$ are proper and lower semicontinuous functions.

Proof: From their definitions, we can verify that they are proper functions. On the other hand, since $\|\cdot\|_*$ is a norm, it is continuous and hence lower semicontinuous. $\delta_{M_R}(\cdot)$ is also lower semicontinuous due to the fact that the $\delta_{M_R}(\cdot)$ is an indicator function of a closed set M_R [38]. ■

Since the sum of lower semicontinuous functions is also lower semicontinuous, we see that Φ in (2) and (4) are both lower semicontinuous.

The lower semicontinuity plays an important role in deriving the proximal maps. For a proper, lower semicontinuous and possibly nonconvex function $g : \mathbb{R}^n \to \mathbb{R}$ and parameter $\tau > 0$, the proximal map $\text{prox}_{g,\tau}$ and Moreau envelope $m_{g,\tau}$ [37] are defined by

$$\text{prox}_{g,\tau}(\mathbf{x}) := \text{argmin}\left\{ g(\mathbf{u}) + \frac{1}{2\tau}\|\mathbf{x} - \mathbf{u}\|_F^2 | \mathbf{u} \in \mathbb{R}^n \right\}$$
$$m_{g,\tau}(\mathbf{x}) := \inf\left\{ g(\mathbf{u}) + \frac{1}{2\tau}\|\mathbf{x} - \mathbf{u}\|_F^2 | \mathbf{u} \in \mathbb{R}^n \right\}.$$

The nonconvexity of g implies that $\text{prox}_{g,\tau}(\mathbf{x})$ may be a set-valued mapping.

The above definitions together with Proposition 1 give the proximal maps for $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$. $\text{prox}_{\|\cdot\|_*,\tau}(\mathbf{X})$ is the matrix shrinkage/soft thresholding (ST) operator [39], [40]

$$\text{prox}_{\|\cdot\|_*,\lambda}(\mathbf{X}) = \mathbf{U}\text{diag}(\max\{\sigma_i - \tau, 0\})\mathbf{V}^T$$

where $\mathbf{X} = \mathbf{U}\text{diag}(\{\sigma_i\}_{1\le i\le r})\mathbf{V}^T$ is the SVD of $\mathbf{X}$. Correspondingly, $\text{prox}_{\delta_{M_R},\tau}(\mathbf{X})$ is the best rank-R approximation to $\mathbf{X}$, and is also called the matrix hard thresholding (HT) operator [36]. Noticing that the parameter τ does not effect $\text{prox}_{\delta_{M_R},\tau}(\mathbf{X})$, we can write it by $\text{prox}_{\delta_{M_R}}(\mathbf{X})$ for short.

B. Algorithm

At first glance, (2) and (4) have a separable structure with respect to $\mathcal{X}_j, 1 \le j \le N$, which can be solved via the conventional block-coordinated descent (BCD) method. BCD solves successively for $j = 1, \ldots, N$ the following subproblems:

$$\mathcal{X}_j^{(k+1)} \in \arg\min_{\mathcal{Z}\in\mathbb{T}} \Phi\big(\mathcal{X}_1^{(k+1)}, \ldots, \mathcal{X}_{j-1}^{(k+1)}, \mathcal{Z}, \mathcal{X}_{j+1}^{(k)}, \mathcal{X}_N^{(k)}\big).$$

However, the above subproblems do not admit closed-form solutions. In view of this, and noticing that J_σ is differentiable everywhere, we can use the linearized and proximal algorithm to approximately solve the subproblems. At the $(k+1)$-iteration, suppose we already have $\mathcal{X}_1^{(k+1)}, \ldots, \mathcal{X}_{j-1}^{(k+1)}$. To obtain $\mathcal{X}_j^{(k+1)}$, we can solve the following subproblem:

$$\mathcal{X}_j^{(k+1)} \in \arg\min_{\mathcal{X}\in\mathbb{T}} \; g(\mathbf{X}_{(j)}) + \frac{\alpha}{2}\big\|\mathcal{X} - \mathcal{X}_j^{(k)}\big\|_F^2 + \left\langle \nabla_{\mathcal{X}_j} J_\sigma\left(\sum_{i=1}^{j-1}\mathcal{X}_i^{(k+1)} + \sum_{i=j}^{N}\mathcal{X}_i^{(k)}\right), \quad \mathcal{X} - \mathcal{X}_j^{(k)}\right\rangle. \tag{7}$$

Here, $\alpha > 0$ is a parameter and $\nabla_{\mathcal{X}_j} J_\sigma(\sum(\boldsymbol{\mathcal{X}}))$ is the partial derivative of J_σ with respect to $\mathcal{X}_j$ that will be specified later. By denoting

$$\mathcal{Y}_j^{(k+1)} = \mathcal{X}_j^{(k)} - \frac{1}{\alpha}\nabla_{\mathcal{X}_j} J_\sigma\left(\sum_{i=1}^{j-1}\mathcal{X}_i^{(k+1)} + \sum_{i=j}^{N}\mathcal{X}_i^{(k)}\right). \tag{8}$$

Equation (7) can be concisely written as

$$\mathcal{X}_j^{(k+1)} \in \arg\min_{\mathcal{X}\in\mathbb{T}} \; \frac{1}{\alpha} g(\mathbf{X}_{(j)}) + \frac{1}{2}\big\|\mathcal{X} - \mathcal{Y}_j^{(k+1)}\big\|_F^2 \tag{9}$$

where by the definition of $g(\cdot)$, we have

$$\mathcal{X}_j^{(k+1)} \in \begin{cases} \text{prox}_{\|\cdot\|_*,\lambda/\alpha}\big(\mathcal{Y}_j^{(k+1)}\big), & \text{if } g(\cdot) = \|\cdot\|_* \\ \text{prox}_{\delta_{M_{R_j}}}\big(\mathcal{Y}_j^{(k+1)}\big), & \text{if } g(\cdot) = \delta_{M_{R_j}}(\cdot). \end{cases} \tag{10}$$

Now, we come to our algorithm. Given an initial guess $\boldsymbol{\mathcal{X}}^{(0)} = \{\mathcal{X}_1^{(0)}, \ldots, \mathcal{X}_N^{(0)}\} \in \mathbb{T}^N$, at each iteration, we compute successively $\mathcal{X}_1^{(k+1)}, \ldots, \mathcal{X}_N^{(k+1)}$ via solving (9). The algorithm stops whenever some stopping criterion is satisfied. This procedure is summarized in Algorithm 1.

Remark 2: To ensure the convergence of the algorithm, a suitable step-size α^{-1} in (8) is preferred. A typical choice is $\alpha > L$, where L is the Lipschitz constant of partial derivative of J_σ, which will be given explicitly later. One can also compute α^{-1} by certain line-search rule, such as the Armijo search rule [41]. Note that a Jacobi update rule can also be employed. That is, in (8), we only use $\mathcal{X}_1^{(k)}, \ldots, \mathcal{X}_N^{(k)}$ to compute $\mathcal{Y}_j^{(k+1)}$ for each j. The advantage is that in each step, the computation of $\mathcal{X}_j^{(k+1)}$ can be parallel. However, an evident drawback is that it cannot utilize the latest information. The stopping criterion can be $\|\boldsymbol{\mathcal{X}}^{(k+1)} - \boldsymbol{\mathcal{X}}^{(k)}\|_F \le \epsilon$, where ϵ is a prescribed number. This will be explained in Section V.

Algorithm 1 Proximal and Linearized BCD With Gauss–Seidel Update Rule (PLiBCD-GS) for (2) and (4)

Input: linear operator $\mathcal{A} : \mathbb{T} \to \mathbb{R}^p$, initial guess $\boldsymbol{\mathcal{X}}^{(0)} \in \mathbb{T}^N$, parameters $R_j, 1 \le j \le N, \sigma > 0, \lambda > 0, \alpha > 0$.
Output: the recovered tensors $\boldsymbol{\mathcal{X}}^{(k+1)} = \{\mathcal{X}_1^{(k+1)}, \ldots, \mathcal{X}_N^{(k+1)}\}$.
while certain stopping criterion is not satisfied **do**
  for $j = 1$ **to** N **do**
    Compute $\mathcal{Y}_j^{(k+1)}$ via the gradient descent step (8).
    Compute $\mathcal{X}_j^{(k+1)}$ via the soft/hard thresholding (10).
  end for
  Set $k := k + 1$.
end while

1) Computing the Gradients: Now, we specify the partial derivative of J_σ with respective to $\mathcal{X}_j$. At $\hat{\boldsymbol{\mathcal{X}}} = \{\hat{\mathcal{X}}_1, \ldots, \hat{\mathcal{X}}_N\}$, the partial derivative $\nabla_{\mathcal{X}_j} J_\sigma(\sum(\hat{\boldsymbol{\mathcal{X}}}))$ associated with the Cauchy loss is given by

$$\sum_{i=1}^{p}(1 + (\langle\mathcal{A}_i, \sum(\hat{\boldsymbol{\mathcal{X}}})\rangle - \mathbf{b}_i)^2/\sigma^2)^{-1}(\langle\mathcal{A}_i, \sum(\hat{\boldsymbol{\mathcal{X}}})\rangle - \mathbf{b}_i)\mathcal{A}_i$$

which can also be further formulated into the following concise form:

$$\nabla_{\mathcal{X}_j} J_\sigma(\sum(\hat{\boldsymbol{\mathcal{X}}})) = \boldsymbol{\mathcal{A}}^\star\Lambda(\boldsymbol{\mathcal{A}}(\sum(\hat{\boldsymbol{\mathcal{X}}})) - \mathbf{b}) \tag{11}$$

where $\Lambda \in \mathbb{R}^{p\times p}$ is a diagonal matrix, with its i-th diagonal entry $\Lambda_{ii} = (1 + (\langle\mathcal{A}_i, \sum(\hat{\boldsymbol{\mathcal{X}}})\rangle - \mathbf{b}_i)^2/\sigma^2)^{-1}$, and $\boldsymbol{\mathcal{A}}^\star$ denotes the adjoint of $\boldsymbol{\mathcal{A}}$. It turns out that for tensor completion problems, the partial derivative enjoys a simpler form

$$\nabla_{\mathcal{X}_j} J_\sigma(\sum(\hat{\boldsymbol{\mathcal{X}}})) = \mathcal{W} * \Lambda * (\sum(\hat{\boldsymbol{\mathcal{X}}}) - \mathcal{B}) \tag{12}$$

where $*$ denotes the Hadamard operator, i.e., the entry-wise product, $\Lambda \in \mathbb{T}$ with $\Lambda_{i_1\ldots i_N} = (1 + (\sum(\hat{\boldsymbol{\mathcal{X}}})_{i_1\ldots i_N} - \mathcal{B}_{i_1\ldots i_N})^2/\sigma^2)^{-1}$, and $\mathcal{W} \in \mathbb{T}$ denotes the mask tensor, i.e., $\mathcal{W}_{i_1\ldots i_N} = 1$ if $(i_1 \ldots i_N) \in \boldsymbol{\Omega}$ and $\mathcal{W}_{i_1\ldots i_N} = 0$ otherwise. Similarly, for the Welsch loss, $\nabla_{\mathcal{X}_j} J_\sigma(\sum(\hat{\boldsymbol{\mathcal{X}}}))$ takes the following form:

$$\sum_{i=1}^{p}\exp(-(\langle\mathcal{A}_i, \sum(\hat{\boldsymbol{\mathcal{X}}})\rangle - \mathbf{b}_i)^2/\sigma^2)(\langle\mathcal{A}_i, \sum(\hat{\boldsymbol{\mathcal{X}}})\rangle - \mathbf{b}_i)\mathcal{A}_i.$$

Remark 3: The above gradients are similar with the one obtained using the least squares estimator, except with the introduction of a weight matrix Λ. We also observe that Λ_{ii} gives a small weight if the corresponding error goes large, which is essentially the key to reduce the influence of outliers.

V. Convergence Results

In this section, we verify the convergence results of proximal and linearized BCD with Gauss–Seidel (PLiBCD-GS) (Algorithm 1) step by step. First, we show that their associated partial derivative is Lipschitz continuous, based on which the descent property of the algorithms can be derived. Then, the convergence results are strengthened by verifying that the function Φ is essentially a Kurdyka–Łojasiewicz (KL) function and using the results of [30].

A. Lipschitz Continuity of the Gradients

First, we show the Lipschitz continuity of the partial derivative $\nabla_{\mathcal{X}_j} J_\sigma(\sum(\boldsymbol{\mathcal{X}}))$, where $\nabla_{\mathcal{X}_j}$ refers to the jth partial derivative, $1 \le j \le N$. Denote $L = \|\boldsymbol{\mathcal{A}}\|_2^2$, where $\|\boldsymbol{\mathcal{A}}\|_2$ is the spectral norm of $\boldsymbol{\mathcal{A}}$.

Proposition 2: For each $j = 1, \ldots, N$ and for any $\mathcal{X}_j, \mathcal{X}_j^+ \in \mathbb{T}$, it holds that

$$\left\| \nabla_{\mathcal{X}_j} J_\sigma \left(\sum_{i \neq j}^{N} \mathcal{X}_i + \mathcal{X}_j^+ \right) - \nabla_{\mathcal{X}_j} J_\sigma \left(\sum_{i=1}^{N} \mathcal{X}_i \right) \right\|_F \le L \left\| \mathcal{X}_j^+ - \mathcal{X}_j \right\|_F.$$

The proof is left to Appendix A.

B. Descent Property

Based on the above propositions, we can derive the descent property for the developed algorithms.

Proposition 3 (Descent Property): For PLiBCD-GS, if the step-size α^{-1} satisfies $\alpha > L$, then it holds that

$$\Phi(\boldsymbol{\mathcal{X}}^{(k+1)}) \le \Phi(\boldsymbol{\mathcal{X}}^{(k)}) - \frac{\alpha - L}{2} \|\boldsymbol{\mathcal{X}}^{(k+1)} - \boldsymbol{\mathcal{X}}^{(k)}\|_F^2. \quad (13)$$

The proof is left to Appendix A.

The above property shows that $\Phi(\boldsymbol{\mathcal{X}}^{(k)})$ is a nonincreasing sequence. Since $\Phi(\cdot)$ is lower bounded by zero, we can deduce that $\|\boldsymbol{\mathcal{X}}^{(k+1)} - \boldsymbol{\mathcal{X}}^{(k)}\|_F \to 0$. Thus, $\|\boldsymbol{\mathcal{X}}^{(k+1)} - \boldsymbol{\mathcal{X}}^{(k)}\|_F \le \epsilon$ can serve as a stopping criterion where ϵ is a prescribed number.

C. Global Convergence

By verifying that Φ satisfies the assumptions in [30], we have Theorem 1.

Theorem 1 (Global Convergence): For PLiBCD-GS, assume that the step-size α^{-1} satisfies $\alpha > L$.

1) If $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ is a sequence generated by PLiBCD-GS for solving the nuclear norm regularized problem (2), then $\sum_{k=1}^{\infty} \|\boldsymbol{\mathcal{X}}^{(k+1)} - \boldsymbol{\mathcal{X}}^{(k)}\|_F < \infty$, and $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ converges to a critical point of problem (2).
2) If $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ is a *bounded* sequence generated by PLiBCD-GS for solving the rank-constrained problem (4), then $\sum_{k=1}^{\infty} \|\boldsymbol{\mathcal{X}}^{(k+1)} - \boldsymbol{\mathcal{X}}^{(k)}\|_F < \infty$, and $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ converges to a critical point of problem (4).

Detailed proofs of this theorem are left to Appendix B. Note that the boundedness of $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ is crucial in the above theorem. The objective function of (2) is coercive due to the sum of nuclear norms term, which together with (13) implies $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ is bounded when solving (2). On the other hand, the objective function of (4) is not coercive, because the nullspace of $\boldsymbol{\mathcal{A}}$ may be nontrivial, and the indicator function of the rank constraint is bounded if the variable does not violate the rank constraint. Therefore, when solving (4), the boundedness of $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ has to be assumed.

Remark 4: We follow [30] in using the term global convergence. In fact, it does not mean that the sequence converges to a global optimizer of the optimization problem. Instead, it states that the whole sequence converges to a critical point, which strengthens the conventional convergence results that every limit point of the sequence is a critical point.

VI. Numerical Experiments

In this section, we present some numerical experiments on synthetic data as well as real data to illustrate the effectiveness of our methods for tensor completion problems. All the numerical computations are conducted on an Intel i7-3770 CPU desktop computer with 16 GB of RAM. The supporting software is MATLAB R2013a. The MATLAB toolbox Tensorlab [42] is employed to perform tensor computations. The following methods are mainly compared.

1) *Our Proposed Method PLiBCD-GS (W, ST):* The PLiBCD-GS method is used to solve (2) with the Welsch loss. ST is short for soft thresholding. This method is abbreviated as W-ST without any confusion. Similar notation C-ST is short for solving (2) with the Cauchy loss.
2) *Our Proposed Method PLiBCD-GS (W, HT):* The PLiBCD-GS method is used to solve (4) with the Welsch loss. HT is short for hard thresholding. W-HT is short for PLiBCD-GS (W, HT). Similar notation C-HT is short for solving (4) with the Cauchy loss.
3) *Four LAD Loss-Based Methods [20]:* S-ST, S-HT, M-ST, and M-HT. They are defined as follows. Denote $L(\mathcal{X}) := \sum_{(i_1 \ldots i_N) \in \Omega} |\mathcal{X}_{i_1 \ldots i_N} - \mathcal{B}_{i_1 \ldots i_N}|$ and $L(\sum(\boldsymbol{\mathcal{X}})) := \sum_{(i_1 \ldots i_N) \in \Omega} |\sum(\boldsymbol{\mathcal{X}})_{i_1 \ldots i_N} - \mathcal{B}_{i_1 \ldots i_N}|$. Then, S-HT and M-ST, respectively, refer to the singleton models

$$\min_{\mathcal{X} \in \mathbb{T}} L(\mathcal{X}) + \lambda \sum_{j=1}^{N} \|\mathbf{X}_{(j)}\|_*$$

and

$$\min_{\mathcal{X} \in \mathbb{T}} L(\mathcal{X}) + \sum_{j=1}^{N} \delta_{M_{R_j}}(\mathbf{X}_{(j)}).$$

M-ST and M-HT, respectively, refer to the mixture models

$$\min_{\mathcal{X} \in \mathbb{T}^N} L(\sum(\boldsymbol{\mathcal{X}})) + \lambda \sum_{j=1}^{N} \|\mathbf{X}_{j,(j)}\|_*$$

and

$$\min_{\mathcal{X} \in \mathbb{T}^N} L(\sum(\boldsymbol{\mathcal{X}})) + \sum_{j=1}^{N} \delta_{M_{R_j}}(\mathbf{X}_{j,(j)}).$$

TABLE I
NOTATIONS IN THE EXPERIMENTS

Notation	Description
ρ_r	the ratio of the rank to the dimensionality of a mode
ρ_o	the ratio of outliers to the number of entries of a tensor
ρ_m	the ratio of missing entries
s	the factor of scale of outliers

4) *The Least Squares Loss-Based Method With Nuclear Norm Heuristic [6] LS-ST:*

$$\min_{\mathcal{X}\in\mathbb{T}} \lambda/2 \sum_{(i_1\ldots i_N)\in\Omega} (\mathcal{X}_{i_1\ldots i_N} - \mathcal{B}_{i_1\ldots i_N})^2 + \sum_{j=1}^{N} \|\mathbf{X}_{(j)}\|_*.$$

Some frequently used notations are introduced in Table I.

A. Synthetic Data

We first generate tensors of size $(50, 50, 50)$ with entries drawn independent identically distributed. From normal distribution, and then truncate the rank of each mode to yield tensors of the Tucker rank $(50\rho_r, 50\rho_r, 50\rho_r)$. The low-rank tensors are denoted as $\mathcal{T}$. The error tensor is denoted by $\mathcal{E}$, with ρ_o entries drawn from some given distributions, while other $(1-\rho_o)$ entries being zero. Then, $\rho_m \times 100\%$ of the entries of $\mathcal{T} + s\mathcal{E}$ is randomly missing. The following three situations are considered.

1) We fix missing ratio $\rho_m = 0.3$, rank ratio $\rho_r = 0.1$, scale factor $s = 1$, and vary outliers ratio ρ_o from 0 to 1.The outliers are drawn from the chi-square distribution.
2) We fix $\rho_r = 0.1$, $\rho_o = 0.2$, $s = 2$, and vary ρ_m from 0 to 1. Outliers are drawn from the chi-square distribution.
3) We fix $\rho_m = 0.3$, $\rho_o = 0.2$, $s = 2$, and vary ρ_r from 0.05 to 1.

The nine methods mentioned in the beginning of this section are tested. The scale parameters σ of the Welsch loss and the Cauchy loss are empirically set to 0.1 and 0.05, respectively. The regularization parameters λ of W-ST and C-ST are set to $0.05n_{\min}/(n_{\max})^{1/2}$. The parameters of S-ST and M-ST are set following the suggestions in [20]. The ranks for the HT type methods are set to the Tucker ranks of $\mathcal{T}$. The initial guess for all the methods is the zero tensor. It seems that since our methods are nonconvex, the results may be potentially influenced by the initial value. However, we have tried different initial guesses and find that the zero tensor gives better or comparable results. Thus for simplicity and reproducibility, we choose zero as our initial guess. The stopping criterion is $\|\mathcal{X}^{(k+1)} - \mathcal{X}^{(k)}\|_F \leq 10^{-4}$ or the iterations exceed 200. The relative error $\|\mathcal{T} - \mathcal{X}\|_F / \|\mathcal{T}\|_F$ is used to measure the recovery ability of the methods. If the error is larger than 1, then we set it to 1. All the results are averaged by ten instances.

Recovery results of the four situations in terms of the relative error are reported in Figs. 2–4, respectively. In these figures, the plots of W-ST and W-HT are in red; the plots of C-ST and C-HT are in green; the plots of S-ST and S-HT are in blue; the plots of M-ST and M-HT are in black; and the plots of LS-ST are in yellow. The plots of the

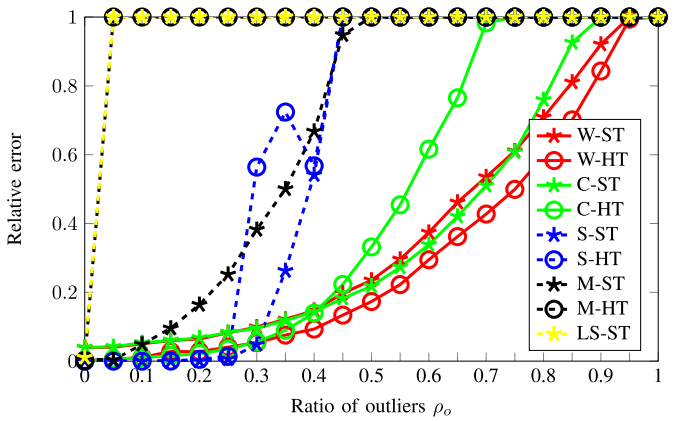


Fig. 2. Recovery results of the third-order tensors with $\rho_m = 0.3$, $\rho_r = 0.1$, $s = 1$ and varying ρ_o. The results show that our methods are more resistant to outliers.

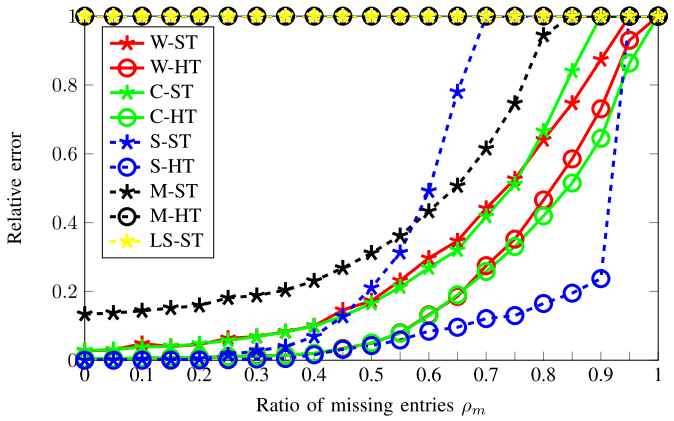


Fig. 3. Recovery results of the third-order tensors with $\rho_o = 0.2$, $\rho_r = 0.1$, $s = 2$ and varying ρ_m, using different methods.

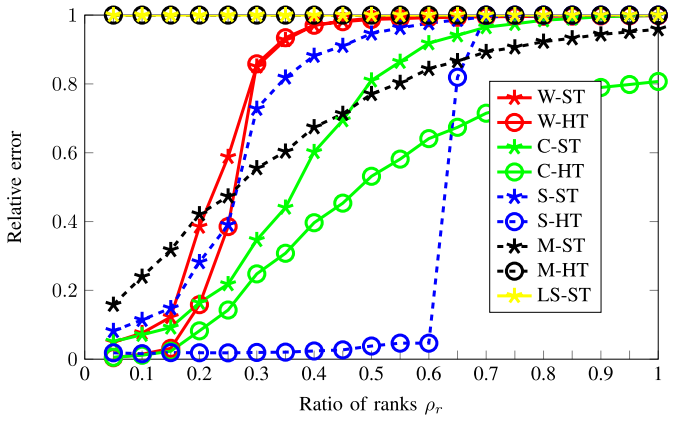


Fig. 4. Recovery results of the third-order tensors with $\rho_m = 0.3$, $\rho_o = 0.2$, $s = 2$, and varying ρ_r, using different methods.

ST type methods are marked with circles $\circ$, while the HT type methods with five-point star $\star$. From all the plots, we observe that LS-ST performs poorly when there exist outliers. M-HT also gives poor results. In fact, the relative error of M-HT is larger than one in almost all the tests. The reason may be that the convergence of the algorithm is not guaranteed [20]. From Fig. 2, we observe that our methods have significantly

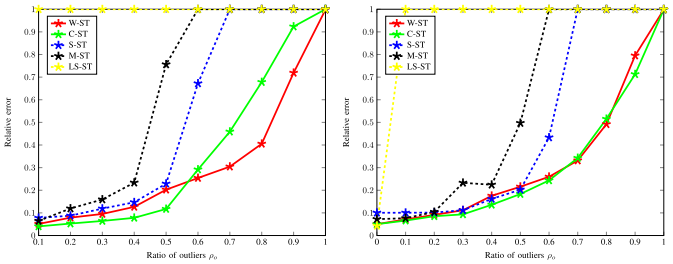
higher recovery accuracy than those in [20]. From Fig. 3, we see that when the level of missing entries increases, our ST type methods perform more stable when $\rho_m > 0.45$. Comparing with the HT type methods, S-HT performs better when $\rho_m > 0.6$. This might be due to that when the true ranks are known, the Tucker model used in S-HT is more suitable to approximate the original tensor. Considering the cases of varying ρ_r and the ST type methods, Fig. 4 shows that C-ST outperforms other methods when $\rho_r \leq 0.45$; for HT type methods, S-HT gets a higher recovery accuracy when $\rho_r \leq 0.6$, which might also be due to the same reason, but the accuracy decreases sharply when $\rho_r > 0.6$. In summary, our methods are more robust in most situations.

B. Traffic Flow Data

In intelligent transportation systems, traffic data, such as traffic flow volume, are collected to support applications in traffic prediction, optimal route choice, and others. Due to some hardware or software problems, the data might be contaminated by outliers, or having missing values. Therefore, some techniques should be performed to recover the original data. The data used here record traffic flow volume from a detector in San Diego and can be downloaded from http://pems.dot.ca.gov. The data consist of 8928 records, each of which represents the traffic flow volume per 5 min, and are collected from January 1–31, 2014. The data can be mapped into the third-order tensor of size $12 \times 24 \times 31$, with modes indicating samples per hour, hours per day, and days per month, respectively. The tensor may have a low-rank structure due to the periodicity. The data can also be modeled as the fourth-order tensor by imposing a week mode since a week trend is also available. In this case, the size of the tensor is $12 \times 24 \times 7 \times 5$, where the data of the last four days of the last week are missing.

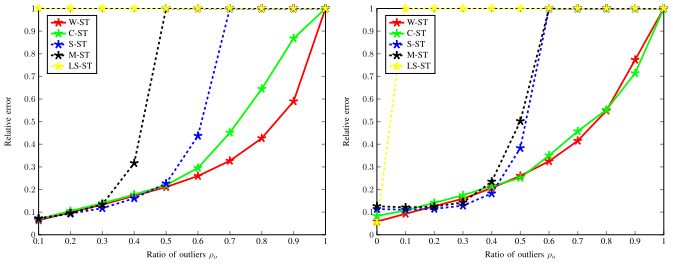
In order to assess the performance of our approaches, randomly chosen entries are corrupted by outliers, with scale being the largest magnitude of the entries. 0% or 20% of the entries are randomly missing. W-ST and C-ST are applied to this problem and compared with S-ST, M-ST, and LS-ST. The scale parameters σ of the Welsch loss and the Cauchy loss are empirically set to 1 and 0.2, respectively. The regularization parameters of our methods are both set to $\lambda = 0.1 n_{\min}/(n_{\max})^{1/2}$. The parameters of S-ST and M-ST are set according to the suggestions in [20]. The initial guesses for all the methods are the zero tensors. The algorithms stop whenever $\|\mathcal{X}^{(k+1)} - \mathcal{X}^{(k)}\|_F \leq 10^{-4}$ or the iterations exceed 500. All the results are averaged over 10 instances.

Relative error results of the algorithms evaluated on the generated third-order and fourth-order tensors are plotted in Figs. 5 and 6, respectively. From the figures, we again observe that LS-ST cannot be resistant with outliers. We can also see that our methods perform better than S-ST and M-ST in most situations. As ρ_o increases, the performances of S-ST and M-ST decrease sharply. Comparing between W-ST and C-ST, it seems that they performs similar when ρ_0 is small, while W-ST is better when ρ_0 increases. This indicates that W-ST is more robust. All the methods seem to have slightly better performance on the third-order tensor model



(a) (b)

Fig. 5. Comparison between W-ST, C-ST, S-ST, M-ST, and LS-ST on traffic data corrupted by outliers and missing entries. The data are represented by the third-order tensor of size $12 \times 24 \times 31$. (a) Fully observed and varying outliers. (b) 20% missing entries and varying outliers.



(a) (b)

Fig. 6. Comparison between W-ST, C-ST, S-ST, M-ST, and LS-ST on traffic data corrupted by outliers and missing entries. The data are represented by the fourth-order tensor of size $12 \times 24 \times 7 \times 5$. (a) Fully observed and varying outliers. (b) 20% missing entries and varying outliers.

than on the fourth-order tensor model when $\rho_0 \leq 0.3$. The reason may be that the data of the last four days of the last week are missing in the fourth-order tensor model, which weakens the low-rank property of the tensor. Finally, we report that in the third-order tensors case, the ranks of the unfoldings of the three recovered factor tensors by our methods are $(\text{rank}(\mathbf{X}_{1,(1)}), \text{rank}(\mathbf{X}_{2,(2)}), \text{rank}(\mathbf{X}_{3,(3)})) = (12, 4, 1)$, respectively, which indicates that the original tensor can be approximated by the sum of three low-rank tensors.

To intuitively illustrate the recovered results, we plot the first 1000 records and the recovered data of C-ST and S-ST in Fig. 7. In each subfigure, the black line represents the raw data, while the red one denotes the recovered data. We can observe that S-ST [Fig. 7(b)] does not perform well in the intervals near 0, 500, and 800, which may be due to the reason that S-ST treats these records as outliers. On the contrary, the data recovered by our method C-ST fit the raw data better, as we observe from Fig. 7(a).

C. Image Inpainting

A color image of size $m \times n$ can be represented as the third-order tensor of size $m \times n \times 3$ by considering the three channels (red, green, and blue) of the image. In real life, images may be corrupted by outliers or partially masked by text. To recover the damaged information, robust tensor completion approaches can be applied. By noticing that the data tensors may be only low rank in the first two modes, the tensor models in our approaches (2) and (4) only consist of two factors, which are low rank in the first two modes, respectively.

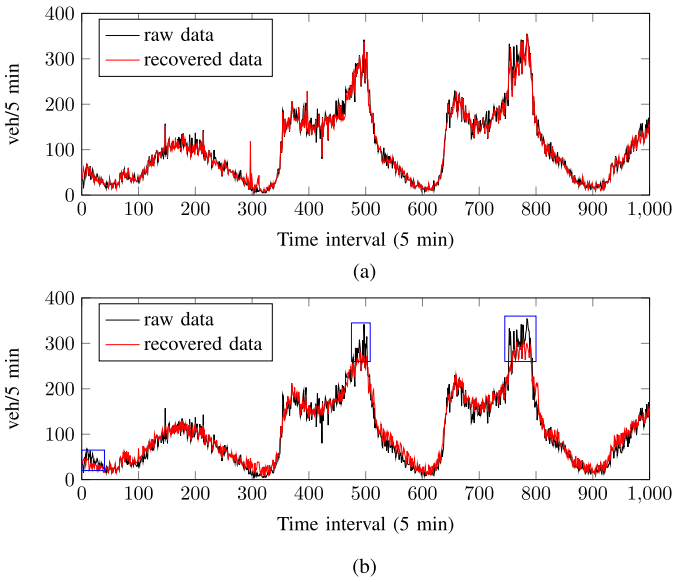


Fig. 7. Plots of the first 1000 records and the recovery results of (a) C-ST and (b) S-ST. The data are corrupted by 30% outliers and 20% missing entries.

The experiments are conducted as follows.

1) Four color images named *House* (256 × 256), *Lena* (225 × 225), *Lake* (512 × 512), and *Pepper* (512 × 512) are tested. Randomly chosen entries are corrupted by outliers with magnitude restricted in $[-1, 1]$. The ratio of outliers ρ_o varies from 0 to 0.4. The images are then covered by some text.
2) The eight robust methods mentioned in the beginning of this section are compared with this experiment. The scale parameters σ of W-ST and C-ST are set to 0.3 and 0.1, respectively, with the regularization parameters being $\lambda = 0.05\min\{m, n\}/(\max\{m, n\})^{1/2}$. To obtain the ranks for W-HT and C-HT, we use the ranks of the factor tensors generated by W-ST and C-ST and rescale them by 0.15, respectively. The parameters of other methods are set following the suggestions in [20]. In accordance with our methods, the tensor models of M-ST and M-HT also contain two factors only. We also compared our method with two of the state-of-the-art methods in [43] and [44], where the method in [43] is based on tensor decomposition and uses prior information of decomposition factors, and is denoted by simultaneous tensor decomposition and completion (STDC)[1]; the second method is based on the PDE, Cahn–Hilliard equation, and is denoted as CH.
3) The initial guesses for all the methods are the zero tensors. The stopping criterion is $\|\mathcal{X}^{(k+1)} - \mathcal{X}^{(k)}\|_F \leq 10^{-4}$, or the iterations exceed 200. The peak-signal-to-noise-ratio (PSNR) value is employed to measure the recovery results.

[1]The MATLAB code was downloaded from https://sites.google.com/site/fallcolor/Tensor%20Completion.rar?attredirects=0, and the parameters of STDC are set the same as those used in the image demo in this code.

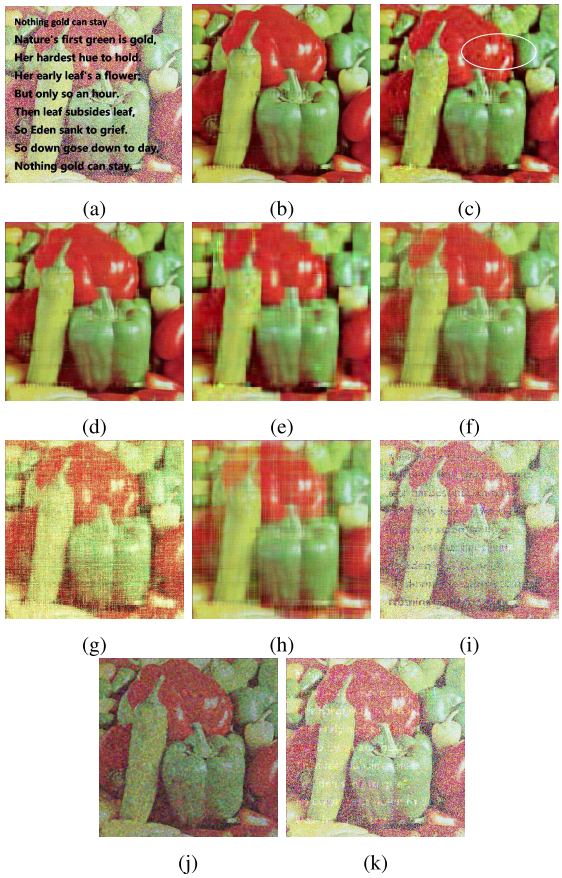


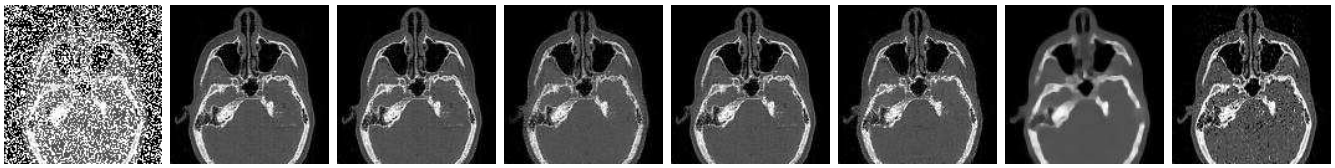
Fig. 8. Recovery results of image *Pepper* contaminated by 40% outliers. (a) Masked. (b) W-ST. (c) W-HT. (d) C-ST. (e) C-HT. (f) S-ST. (g) S-HT. (h) M-ST. (i) M-HT. (j) CH. (k) STDC.

The PSNR values of images recovered by different methods are reported in Table II. From the table, we see that when there are no outliers, the STDC method outperforms other methods. When there exist outliers, the performance of STDC decreases sharply. In this setting, our method W-ST achieves the highest values among all the images and all the situations, followed by C-ST. This may be because the Welsch loss is more robust than the Cauchy loss, which can be implied from the discussions about the weight of outliers in Section III-C and shown in Fig. 1. LAD loss-based methods are more sensitive to outliers than ours. Considering the methods with the hard thresholding, namely, W-HT, we observe that the HT type methods do not perform as well as the ST type methods, which is due to the reason that the HT type methods cannot completely remove the text from the images. To intuitively illustrate the effectiveness of our methods, recovered images of *Pepper* are presented in Fig. 8. We can see that the image recovered by W-ST can retain more details and remove the text, followed by C-ST. S-ST and M-ST can also remove the text but lose more details than our methods, which may be that they are not as robust as our methods. W-HT also retains more details than other HT type methods, but some text cannot be totally removed, as masked by the ellipsoid in Fig. 8(c). S-HT and M-HT cannot remove the outliers. In fact, we have tried different choices of ranks while get similar results. We also observe that nonrobust methods are influenced by outliers seriously.

TABLE II

PSNR VALUES OF THE RECOVERY IMAGES BY DIFFERENT METHODS. THE ORIGINAL IMAGES ARE CONTAMINATED BY DIFFERENT LEVELS OF OUTLIERS AND MASKED BY TEXT. THE OUTLIERS RATIO ρ_o VARIES FROM 10% TO 40%

Image	ρ_o	W-ST	W-HT	C-ST	C-HT	S-ST [20]	S-HT [20]	M-ST [20]	M-HT [20]	CH [44]	STDC [43]
House	0	29.86	24.25	27.47	22.75	26.22	14.47	27.13	15.70	32.24	**32.93**
	10%	**29.33**	22.78	26.88	22.42	25.10	15.98	26.54	10.59	11.93	10.28
	20%	**28.70**	23.63	26.20	22.24	24.86	12.89	25.45	7.02	13.40	6.99
	30%	**27.90**	24.13	25.31	21.51	24.11	9.51	23.27	4.95	14.98	5.30
	40%	**26.99**	23.20	24.28	19.93	21.48	6.14	19.22	3.56	16.49	4.15
Lake	0	27.78	23.43	26.22	22.61	23.90	15.16	25.22	16.59	27.87	**29.70**
	10%	**27.31**	23.23	25.52	22.42	23.68	15.03	24.27	10.34	12.50	10.40
	20%	**26.72**	23.20	24.67	22.18	23.30	11.33	23.41	6.79	13.79	7.34
	30%	**26.00**	22.85	23.62	21.60	22.10	8.08	20.94	4.82	15.08	5.58
	40%	**25.07**	22.46	22.23	20.63	19.71	5.58	19.48	3.57	15.97	4.33
Lena	0	29.23	25.29	27.88	24.71	26.61	15.30	26.95	16.26	31.07	**31.77**
	10%	**28.83**	25.22	27.33	24.47	25.80	15.40	26.45	10.65	12.49	10.48
	20%	**28.33**	25.08	26.65	23.95	25.41	11.91	25.45	6.99	13.95	7.39
	30%	**27.73**	24.86	25.85	23.14	23.97	8.63	23.11	4.97	15.41	5.62
	40%	**26.97**	24.09	24.73	22.30	21.32	5.74	19.45	3.61	16.68	4.35
Pepper	0	29.67	25.00	28.20	23.33	26.52	16.30	27.99	17.52	32.71	**32.91**
	10%	**29.22**	25.03	27.50	22.87	26.03	15.20	27.06	10.60	13.29	10.42
	20%	**28.67**	24.85	26.66	22.49	25.32	11.21	25.50	6.98	14.67	7.31
	30%	**27.97**	24.69	25.61	21.84	23.48	8.12	21.66	4.97	15.87	5.56
	40%	**27.11**	24.37	24.30	20.86	20.15	5.48	18.28	3.62	16.57	4.32



(a) One slice (b) W-ST (c) C-ST (d) S-ST [20] (e) M-ST [20] (f) LS-ST [6] (g) CH [44] (h) STDC [43]

Fig. 9. Recovery results of one slice of the INCISIX data set. (a) One slice with 40% entries missing. (b)–(h) Recovered results of different methods.

TABLE III

RELATIVE ERRORS OF DIFFERENT METHODS ON RECOVERING THE INCISIX (D1), KNIX1 (D2), AND KNIX2 (D3) DATA SETS

	ρ_m	W-ST	C-ST	S-ST [20]	M-ST [20]	LS-ST [6]	CH [44]	STDC [43]
D1	0.4	**0.1075**	0.1095	0.1489	0.1205	0.1386	0.2078	0.1282
	0.6	0.1765	**0.1738**	0.2738	0.2008	0.2057	0.2398	0.1913
	0.8	0.2940	**0.2906**	0.4495	0.3654	0.3361	0.2917	0.3766
D2	0.4	**0.0948**	0.0960	0.1403	0.1194	0.1009	0.4101	0.1169
	0.6	0.1719	0.1740	0.2005	0.1788	**0.1581**	0.4595	0.1653
	0.8	0.2441	**0.2418**	0.3440	0.2726	0.2511	0.5689	0.2608
D3	0.4	0.0969	0.0979	0.1634	0.1326	0.0930	0.2683	**0.0937**
	0.6	0.1394	0.1382	0.1983	0.1692	0.1635	0.3137	**0.1355**
	0.8	0.2357	**0.2322**	0.3805	0.3066	0.2863	0.3863	0.2486

D. MRI Data

We test different methods on the MRI data set INCISIX ($166 \times 128 \times 128$) and KNIX from http://www.osirix-viewer.com/datasets/. From KNIX, we choose two data sets, each of which is of size $22 \times 128 \times 128$. We try W-ST, C-ST, S-ST, M-ST, LS-ST, CH, and STDC on the data with different ratios of missing entries. The parameters are set the same as the previous experiment. The model LS-ST is solved iteratively with an increasing λ value because of the absence of noise. The relative errors are reported in Table III, along with the recovered results of one slice of INCISIX presented in Fig. 9. The results show that in most cases, W-ST and C-ST perform better than other methods. The reason may be that the raw data can be better approximated by the mixture tensor model. CH performs worst, which may be that the PDE type methods do not explore the low-rank structure of the tensor to be recovered. The results also indicate that our methods are safe when there is no noise or outliers.

E. Removing Shadows and Specularities From Face Images

We test our methods on the YaleB data set. The goal is to remove shadows and specularities from face images. The data set consists of 64 face images of size 192×168, which gives a tensor of size $64 \times 192 \times 168$. The shadows and specularities can be seen as outliers [34]. We apply W-ST ($\sigma = 0.5$, $\lambda = 5 \cdot \min_i\{n_i\}/(\max_i\{n_i\})^{1/2}$) on the data set and the methods can successfully process the images. We compare our methods with the robust PCA (RPCA) approach [34] (recommended parameter $\lambda = 1/(\max\{n_1, n_2 \cdot n_3\})^{1/2}$ [34]), which is solved by the inexact augmented Lagrange multiplier (IALM) algorithm proposed in [45]. Results are reported in Fig. 10. It seems that our method may better remove the shadows and specularities, particularly from those presented in the first row of Fig. 10.

F. Video Surveillance

For the last experiment, we test our methods on the video surveillance data available at http://perception.i2r.a-star.edu.sg/bk_model/bk_index.html. The aim is to detect the moving objects from static background. The moving objects can be regarded as outliers. The data set used here consists of

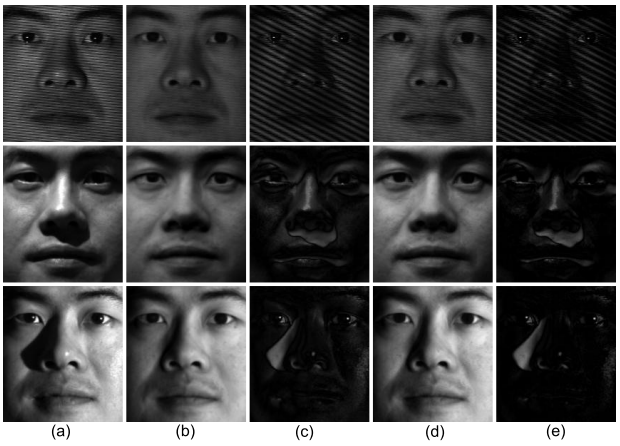


Fig. 10. Results of some face images. (a) Three original images. (b) Faces recovered by W-ST. (c) Shadows and specularities of W-ST. (d) Faces recovered by RPCA. (e) Shadows and specularities of RPCA. The results show that our method may better remove shadows and specularities from faces images.

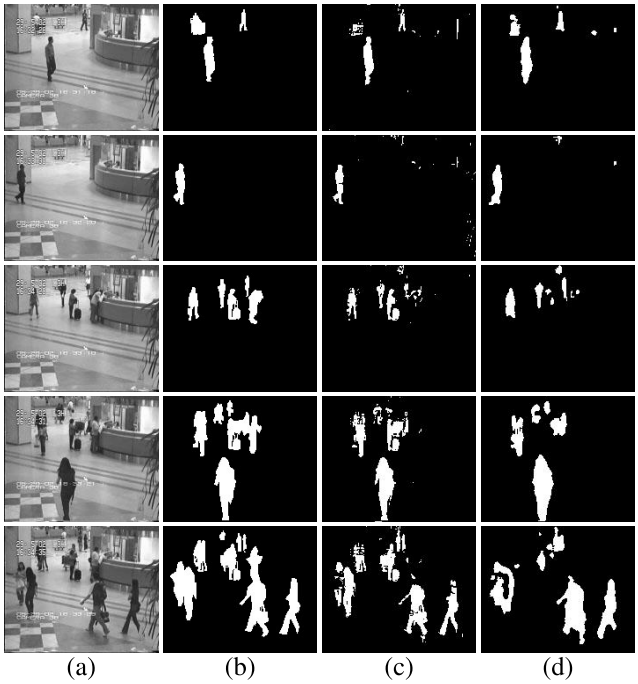


Fig. 11. Extracting results of some frames of the video airport (frame number: 1656, 2289, 2810, 2926, and 2961). (a) Original frames. (b) Ground truth. (c) Extracted by W-ST. (d) Extracted by RPCA.

4583 color frames of size 144×176, from which we choose the last 3583 frames, giving a tensor of size $3583 \times 144 \times 176$, otherwise the algorithms are too time consuming. We compare W-ST ($\sigma = 0.5$, $\lambda = 10 \cdot \min_i\{n_i\}/(\max_i\{n_i\})^{1/2}$) with RPCA (recommended parameter $\lambda = 1/(\min\{n_1, n_2 \cdot n_3\})^{1/2}$, [34]) solved by the IALM algorithm [45]. To quantitatively compare the performance, we report their F-measure value on 20 images with ground truth, where

$$F\text{-measure} = \frac{2 \cdot \text{precision} \cdot \text{recall}}{\text{precision} + \text{recall}}$$

TABLE IV
COMPARISON OF THE F-MEASURE VALUE ON VIDEO AIRPORT. THE SECOND COLUMN SHOWS THE VALUE OVER 20 TEST IMAGES; THE THIRD TO THE LAST COLUMNS SHOW THE VALUES ON SOME SPECIFIC IMAGES

Alg.	Averaged	No. 1656	No. 2289	No. 2810	No. 2926	No. 2961
W-ST	0.6944	0.7651	0.7851	0.6968	0.8060	0.7808
RPCA [45]	0.6730	0.8271	0.8411	0.5593	0.7548	0.7052
	No. 3049	No. 3434	No. 3800	No. 4048	No. 4333	No. 4432
W-ST	0.7296	0.7109	0.4898	0.6036	0.5498	0.5614
RPCA [45]	0.4771	0.7206	0.2783	0.6794	0.5374	0.6176

here, precision $= tp/(tp + fn)$ and recall $= tp/(tp + fp)$, where tp, fp, and fn are the number of pixels of correctly marked foreground, wrongly marked foreground, and wrongly marked background, respectively. The results are reported in Table IV, where the second column shows the F-measure over 20 test images; the third to the last columns show the values on some specific images. The results show that W-ST performs slightly better. Intuitive results are shown in Fig. 11 as well.

VII. CONCLUSION

This paper addressed the robust tensor recovery problems using regularized redescending M-estimators, where the Welsch loss and the Cauchy loss were employed. To solve the nonconvex problems, the linearized and proximal block coordinate methods were developed, and global convergence was verified. The numerical experiments on synthetic data and real data under various circumstances were conducted. Empirically, we find that our models are more robust in the presence of outliers and can give at least comparable performance in the absence of outliers. Thus, it is safe to use our methods in tensor recovery problems under various circumstances.

APPENDIX A
PROOF OF PROPOSITION 2 AND PROPOSITION 3

Proof of Proposition 2: To prove the conclusion, it suffices to show that for any $\mathcal{X}^+, \mathcal{X} \in \mathbb{T}$

$$\|\nabla J_\sigma(\mathcal{X}^+) - \nabla J_\sigma(\mathcal{X})\|_F \le L\|\mathcal{X}^+ - \mathcal{X}\|_F. \tag{14}$$

We first observe that

$$\begin{aligned}\|\nabla J_\sigma(\mathcal{X}^+) - \nabla J_\sigma(\mathcal{X})\|_F &= \|\boldsymbol{\mathcal{A}}^*\Lambda_+(\boldsymbol{\mathcal{A}}(\mathcal{X}^+) - b) - \boldsymbol{\mathcal{A}}^*\Lambda(\boldsymbol{\mathcal{A}}(\mathcal{X}) - b)\|_F \\ &\le \|\boldsymbol{\mathcal{A}}\|_2\|\Lambda_+(\boldsymbol{\mathcal{A}}(\mathcal{X}^+) - b) - \Lambda(\boldsymbol{\mathcal{A}}(\mathcal{X}) - b)\|_F \end{aligned} \tag{15}$$

where Λ_+ and Λ are, respectively, the diagonal matrices corresponding to $\nabla J_\sigma(\mathcal{X}^+)$ and $\nabla J_\sigma(\mathcal{X})$. Let $\mathbf{z}^+ = \boldsymbol{\mathcal{A}}(\mathcal{X}^+) - b$ and $\mathbf{z} = \boldsymbol{\mathcal{A}}(\mathcal{X}) - b$. Associated with the Cauchy loss, we have

$$\begin{aligned}&\|\Lambda_+(\boldsymbol{\mathcal{A}}(\mathcal{X}^+) - b) - \Lambda(\boldsymbol{\mathcal{A}}(\mathcal{X}) - b)\|_F^2 \\ &\quad= \sum_{i=1}^{p}\left(\left(1 + \left(\mathbf{z}_i^+\right)^2/\sigma^2\right)^{-1}\mathbf{z}_i^+ - \left(1 + \mathbf{z}_i^2/\sigma^2\right)^{-1}\mathbf{z}_i\right)^2.\end{aligned}$$

Associated with the Welsch loss, we have

$$\|\Lambda_+(\mathcal{A}(\mathcal{X}^+)-b)-\Lambda(\mathcal{A}(\mathcal{X})-b)\|_F^2 = \sum_{i=1}^{p}\left(\exp\left(-\left(\mathbf{z}_i^+\right)^2/\sigma^2\right)\mathbf{z}_i^+-\exp\left(-\mathbf{z}_i^2/\sigma^2\right)\mathbf{z}_i\right)^2.$$

For the Cauchy loss, it can be verified that the magnitude of the derivative of function $(1+t^2/\sigma^2)^{-1}t$ is not larger than 1 for any $t\in\mathbb{R}$ and $\sigma>0$. Therefore, it follows from the mean value theorem that for any $t_1,t_2\in\mathbb{R}$ and $\sigma>0$:

$$\left|\left(1+t_1^2/\sigma^2\right)^{-1}t_1-\left(1+t_2^2/\sigma^2\right)^{-1}t_2\right|\le|t_1-t_2|.$$

A similar inequality holds for the Welsch loss. Thus, we obtain

$$\|\Lambda_+(\mathcal{A}(\mathcal{X}^+)-b)-\Lambda(\mathcal{A}(\mathcal{X})-b)\|_F \le \|\mathcal{A}\mathcal{X}^+-\mathcal{A}\mathcal{X}\|_F\le\|\mathcal{A}\|_2\|\mathcal{X}^+-\mathcal{X}\|_F.$$

This in connection with (15) implies (14). The proof is completed. ■

Proof of Proposition 3: First, since $\mathcal{X}_j^{(k+1)}$ is an optimal solution of (7), we get

$$\frac{\alpha}{2}\left\|\mathcal{X}_j^{(k+1)}-\mathcal{X}_j^{(k)}\right\|_F^2+\left\langle\nabla_{\mathcal{X}_j}J_\sigma,\mathcal{X}_j^{(k+1)}-\mathcal{X}_j^{(k)}\right\rangle+g\left(X_{j,(j)}^{(k+1)}\right)\le g\left(X_{j,(j)}^{(k)}\right) \tag{16}$$

where $\nabla_{\mathcal{X}_j}J_\sigma$ is short for the partial derivative of J_σ, $\nabla_{\mathcal{X}_j}J_\sigma(\sum_{i=1}^{j-1}\mathcal{X}_i^{(k+1)}+\sum_{i=j}^{N}\mathcal{X}_i^{(k)})$. Then, Proposition 2 implies that

$$J_\sigma\left(\sum_{i=1}^{j}\mathcal{X}_i^{(k+1)}+\sum_{i=j+1}^{N}\mathcal{X}_i^{(k)}\right)\le J_\sigma\left(\sum_{i=1}^{j-1}\mathcal{X}_i^{(k+1)}+\sum_{i=j}^{N}\mathcal{X}_i^{(k)}\right)+\left\langle\nabla_{\mathcal{X}_j}J_\sigma,\mathcal{X}_j^{(k+1)}-\mathcal{X}_j^{(k)}\right\rangle+\frac{L}{2}\left\|\mathcal{X}_j^{(k+1)}-\mathcal{X}_j^{(k)}\right\|_F^2. \tag{17}$$

Combining (16) and (17) yields

$$J_\sigma\left(\sum_{i=1}^{j}\mathcal{X}_i^{(k+1)}+\sum_{i=j+1}^{N}\mathcal{X}_i^{(k)}\right)+g\left(X_{j,(j)}^{(k+1)}\right)\le J_\sigma\left(\sum_{i=1}^{j-1}\mathcal{X}_i^{(k+1)}+\sum_{i=j}^{N}\mathcal{X}_i^{(k)}\right)+g\left(X_{j,(j)}^{(k)}\right)-\frac{a-L}{2}\left\|\mathcal{X}_j^{(k+1)}-\mathcal{X}_j^{(k)}\right\|_F^2. \tag{18}$$

Summing (18) from $j=1$ to N thus gives (13). ■

Appendix B
Kurdyka–Łojasiewicz Property and the Global Convergence

Suppose $\{\mathbf{x}^{(k)}\}$ is a sequence generated by an algorithm applied to a nonconvex optimization problem. Recently, the promising work in [30] ([30, Th. 1] and [30, Sec. 3.6]) shows that, if the objective function of the problem satisfies the KL property and some additional assumptions, and if the proximal and linearized type methods are applied to the problem, then the sequence $\{\mathbf{x}^{(k)}\}$ converges to a critical point.

We first verify that the objective function Φ defined in (5) is exactly a KL function, and then verify that Φ meets assumptions required in [30], hence showing the global convergence of PLiBCD-GS. Here, we do not go into the detail of the KL functions, instead, we only identify the KL property of Φ using some properties provided in [30].

Proposition 4 (See [30]): A function is KL if it is subanalytic or it is both lower semicontinuous and semialgebraic.

We first have the following results.

Proposition 5: The data-fitting risk J_σ is analytic.

Proof: By the additivity of the analytic functions, it suffices to show the Cauchy loss and the Welsch loss are analytic. Because they are, respectively, the composition of logarithm and quadratic function, and composition of exponential and quadratic function, by the properties of the analytic functions, it follows that both of them are analytic functions. ■

Semialgebraic property of $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$:

Proposition 6: Both $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$ are semialgebraic functions.

To investigate the semialgebraic property of $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$, some definitions and properties concerning the semialgebraic sets and functions will be introduced briefly.

Definition 2 (Semialgebraic Sets and Functions, See [46]): A set $C\subset\mathbb{R}^n$ is called semialgebraic if there exists a finite number of real polynomials $p_{ij},q_{ij}:\mathbb{R}^n\to\mathbb{R}$, $1\le i\le s$, $1\le j\le t$ such that

$$C=\bigcup_{i=1}^{s}\bigcap_{j=1}^{t}\{\mathbf{x}\in\mathbb{R}^n|\,p_{ij}(\mathbf{x})=0,\quad q_{ij}(\mathbf{x})>0\}.$$

A function $\phi:\mathbb{R}^n\to\mathbb{R}\setminus\{-\infty\}$ is called semialgebraic if its graph

$$\text{graph }\phi:=\{(\mathbf{x},t)\in\mathbb{R}^{n+1}|\phi(\mathbf{x})=t\}$$

is a semialgebraic set.

Operations that keep the semialgebraic property of sets include finite unions, finite intersections, complementation, and Cartesian products. The following theorem is also helpful for identifying semialgebraic sets.

Theorem 2 (Tarski–Seidenberg Theorem, See [47]): Let C be a semialgebraic set of $\mathbb{R}^n$. Then, the image $\pi(C)$ is a semialgebraic set if the following holds.

1) $\pi:\mathbb{R}^n\to\mathbb{R}^{n-1}$ is the projection on the first $(n-1)$ coordinates.
2) $\pi:\mathbb{R}^n\to\mathbb{R}^m$ is a polynomial mapping.

Some common examples of semialgebraic functions are listed in the following, which are provided in [30].

1) Real polynomial functions.
2) Characteristic functions of semialgebraic sets.
3) Finite sums, product, and composition of semialgebraic functions.
4) Supreme type functions: $\sup\{g(\mathbf{x},\mathbf{y})|\mathbf{y}\in C\}$ is semialgebraic if g and C are semialgebraic.

Based on the above definitions and properties, we are able to prove the semialgebraic property of $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$.

Proof of Proposition 6: We first prove the semialgebraic property of $\|\cdot\|_*$. Recall that for any matrix $\mathbf{X}\in\mathbb{R}^{m_1\times m_2}$, the dual formulation of $\|\mathbf{X}\|_*$ is given by

$$\|\mathbf{X}\|_*=\sup\{\langle\mathbf{X},\mathbf{Y}\rangle\,|\|\mathbf{Y}\|_2\le1\}=\sup\{\langle\mathbf{X},\mathbf{Y}\rangle\,|\|\mathbf{Y}\|_2=1\}.$$

Since $\langle \mathbf{X}, \mathbf{Y} \rangle$ is a polynomial, it suffices to show the level set of the spectral norm $\{\mathbf{Y} \in \mathbb{R}^{m_1 \times m_2} | \|\mathbf{Y}\|_2 = 1\}$ is a semialgebraic set. To prove this, we need to first show that $\|\cdot\|_2$ is a semialgebraic function. This can be verified by adopting the following representation:

$$\|\mathbf{Y}\|_2 = \max \left\{ \mathbf{z}_1^T \mathbf{Y} \mathbf{z}_2 | \mathbf{z}_i \in \mathbb{R}^{m_i}, \quad \|\mathbf{z}_i\| = 1, \ i = 1, 2 \right\}$$

and the semialgebraic property of the set $\{\mathbf{z}_i \in \mathbb{R}^{m_i} | \|\mathbf{z}_1\| = 1\}$, $i = 1, 2$. Since $\|\cdot\|_2$ is semialgebraic, by definition, the graph of the spectral norm

$$\text{graph } \|\cdot\|_2 = \{(\mathbf{Y}, t) | \|\mathbf{Y}\|_2 = t\}$$

is a semialgebraic set. Applying Tarski–Seidenberg theorem to graph$\|\cdot\|_2$ thus yields the semialgebraic property of the level set of $\|\cdot\|_2$. As a result, we have proved $\|\cdot\|_*$ is a semialgebraic function.

We then show the semialgebraic property for $\delta_{M_R}(\cdot)$. This can be verified by showing that the same property holds for set M_R. Define $\varphi : \mathbb{R}^{m_1 \times R} \times \mathbb{R}^{m_2 \times R} \to \mathbb{R}^{m_1 \times m_2}$ by $\varphi(\mathbf{U}, \mathbf{V}) = \mathbf{U}\mathbf{V}^T$. It is clear that φ is a polynomial, and the image of φ is exactly the set M_R. Therefore, the semialgebraic property of M_R follows again from Tarski–Seidenberg theorem. The proof has been completed. ■

Propositions 1, 4, and 6 show that $\|\cdot\|_*$ and $\delta_{M_R}(\cdot)$ are KL functions. This in connection with Proposition 5 and the additivity of KL functions yields the following proposition.

Proposition 7: The function Φ defined in (5) is a KL function.

We then can prove Theorem 1.

Proof of Theorem 1: According to [30, Th. 1] and the discussions in [30, Sec. 3.6], to prove Theorem 1, besides the KL property of Φ, we also need to verify that: 1) the regularization term is proper and lower semicontinuous; 2) the loss term is continuously differentiable; 3) the gradient of the loss term is Lipschitz continuous; and 4) $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ is bounded. Item 1) has been shown in Proposition 1. For 2) and 3), we have for any $\boldsymbol{\mathcal{X}}^+, \boldsymbol{\mathcal{X}} \in \mathbb{T}^N$

$$\begin{aligned}
&\|\nabla J_\sigma(\textstyle\sum(\boldsymbol{\mathcal{X}}^+)) - \nabla J_\sigma(\sum(\boldsymbol{\mathcal{X}}))\|_F^2 \\
&= \sum_{j=1}^{N} \|\nabla_{\mathcal{X}_j} J_\sigma(\textstyle\sum(\boldsymbol{\mathcal{X}}^+)) - \nabla_{\mathcal{X}_j} J_\sigma(\sum(\boldsymbol{\mathcal{X}}))\|_F^2 \\
&\le N\|\boldsymbol{\mathcal{A}}\|_2^4 \|\textstyle\sum(\boldsymbol{\mathcal{X}}^+) - \sum(\boldsymbol{\mathcal{X}})\|_F^2 \le N^2 \|\boldsymbol{\mathcal{A}}\|_2^4 \|\boldsymbol{\mathcal{X}}^+ - \boldsymbol{\mathcal{X}}\|_F^2
\end{aligned}$$

where $\nabla_{\mathcal{X}_j}$ refers to the jth partial derivative, and the first inequality is similar to the proof of Proposition 2 in Appendix A. Therefore, 2) and 3) are verified simultaneously. If $g(\cdot)$ takes the nuclear norm, then the boundedness of $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ follows from the coercivity of $J(\cdot) + \sum \|\cdot\|_*$; if $g(\cdot)$ is the indicator function of the rank constraint, then the boundedness follows from the assumption of Theorem 1, and this gives 4). The above results in connection with Proposition 7, [30, Th. 1], and [30, Sec. 3.6] give Theorem 1. ■

Remark 5: The boundedness of $\{\boldsymbol{\mathcal{X}}^{(k)}\}$ is assumed when $g(\cdot)$ is the indicator function of the rank constraint, because $J(\cdot) + \sum \delta_{M_{R_j}}(\cdot)$ is not coercive due to the nullspace of $\boldsymbol{\mathcal{A}}$ may be nontrivial.

ACKNOWLEDGMENT

This paper reflects only the authors' views, the Union is not liable for any use that may be made of the contained information.

REFERENCES

[1] L. R. Tucker, "Some mathematical notes on three-mode factor analysis," *Psychometrika*, vol. 31, no. 3, pp. 279–311, 1966.

[2] 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, 1970.

[3] R. A. Harshman, "Foundations of the PARAFAC procedure: Models and conditions for an 'explanatory' multi-modal factor analysis," *UCLA Working Papers Phonetics*, vol. 16, pp. 1–84, Dec. 1970.

[4] L. De Lathauwer, B. De Moor, and J. Vandewalle, "A multilinear singular value decomposition," *SIAM J. Matrix Anal. Appl.*, vol. 21, no. 4, pp. 1253–1278, 2000.

[5] J. Liu, P. Musialski, P. Wonka, and J. Ye, "Tensor completion for estimating missing values in visual data," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 208–220, Jan. 2013.

[6] S. Gandy, B. Recht, and I. Yamada, "Tensor completion and low-n-rank tensor recovery via convex optimization," *Inverse Problems*, vol. 27, no. 2, p. 025010, 2011.

[7] N. Li and B. Li, "Tensor completion for on-board compression of hyperspectral images," in *Proc. 17th IEEE Int. Conf. Image Process.*, Hong Kong, Sep. 2010, pp. 517–520.

[8] X. Geng, K. Smith-Miles, Z.-H. Zhou, and L. Wang, "Face image modeling by multilinear subspace analysis with missing values," *IEEE Trans. Syst., Man, Cybern. B, Cybern.*, vol. 41, no. 3, pp. 881–892, Jun. 2011.

[9] M. Signoretto, Q. T. Dinh, L. De Lathauwer, and J. A. K. Suykens, "Learning with tensors: A framework based on convex optimization and spectral regularization," *Mach. Learn.*, vol. 94, no. 3, pp. 303–351, 2014.

[10] M. Signoretto, R. Van de Plas, B. De Moor, and J. A. K. Suykens, "Tensor versus matrix completion: A comparison with application to spectral data," *IEEE Signal Process. Lett.*, vol. 18, no. 7, pp. 403–406, Jul. 2011.

[11] H. Tan, J. Feng, G. Feng, W. Wang, and Y.-J. Zhang, "Traffic volume data outlier recovery via tensor model," *Math. Problems Eng.*, vol. 2013, Feb. 2013, Art. ID 164810.

[12] R. Tomioka, K. Hayashi, and H. Kashima. (Oct. 2010). "Estimation of low-rank tensors via convex optimization." [Online]. Available: http://arxiv.org/abs/1010.0789

[13] R. Tomioka and T. Suzuki, "Convex tensor decomposition via structured Schatten norm regularization," in *Proc. Adv. Neural Inf. Process. Syst.*, Stateline, NV, USA, Dec. 2013, pp. 1331–1339.

[14] B. Huang, C. Mu, D. Goldfarb, and J. Wright. (2014). *Provable Low-Rank Tensor Recovery*. [Online]. Available: http://www.optimization-online.org/DB_HTML/2014/02/4252.html

[15] H. Rauhut, R. Schneider, and Z. Stojanac, "Low rank tensor recovery via iterative hard thresholding," in *Proc. 10th Int. Conf. Sampling Theory Appl.*, Bremen, Germany, Jul. 2013, pp. 21–24.

[16] L. Yang, Z.-H. Huang, and X. Shi, "A fixed point iterative method for low n-rank tensor pursuit," *IEEE Trans. Signal Process.*, vol. 61, no. 11, pp. 2952–2962, Jun. 2013.

[17] X. Zhang, D. Wang, Z. Zhou, and Y. Ma, "Simultaneous rectification and alignment via robust recovery of low-rank tensors," in *Proc. Adv. Neural Inf. Process. Syst.*, Stateline, NV, USA, Dec. 2013, pp. 1637–1645.

[18] Y. Li, J. Yan, Y. Zhou, and J. Yang, "Optimum subspace learning and error correction for tensors," in *Proc. 11th Eur. Conf. Comput. Vis.*, vol. 6383. Crete, Greece, Sep. 2010, pp. 790–803.

[19] H. Tan, B. Cheng, J. Feng, G. Feng, and Y. Zhang, "Tensor recovery via multi-linear augmented Lagrange multiplier method," in *Proc. 6th Int. Conf. Image Graph.*, Hefei, China, Aug. 2011, pp. 141–146.

[20] 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, 2014.

[21] P. J. Huber and E. M. Ronchetti, *Robust Statistics*. New York, NY, USA: Wiley, 2009.

[22] K. Inoue, K. Hara, and K. Urahama, "Robust multilinear principal component analysis," in *Proc. IEEE 12th Int. Conf. Comput. Vis.*, Kyoto, Japan, Sep./Oct. 2009, pp. 591–597.

[23] G. Shevlyakov, S. Morgenthaler, and A. Shurygin, "Redescending M-estimators," *J. Statist. Planning Inference*, vol. 138, no. 10, pp. 2906–2917, 2008.

[24] I. Mizera and C. H. Müller, "Breakdown points of Cauchy regression-scale estimators," *Statist. Probab. Lett.*, vol. 57, no. 1, pp. 79–89, 2002.

[25] X. Wang, Y. Jiang, M. Huang, and H. Zhang, "Robust variable selection with exponential squared loss," *J. Amer. Statist. Assoc.*, vol. 108, no. 502, pp. 632–643, 2013.

[26] Y. Feng, X. Huang, L. Shi, Y. Yang, and J. A. K. Suykens, "Learning with the maximum correntropy criterion induced losses for regression," *J. Mach. Learn. Res.*, in press.

[27] W. Liu, P. P. Pokharel, and J. C. Principe, "Correntropy: Properties and applications in non-Gaussian signal processing," *IEEE Trans. Signal Process.*, vol. 55, no. 11, pp. 5286–5298, Nov. 2007.

[28] R. He, W.-S. Zheng, and B.-G. Hu, "Maximum correntropy criterion for robust face recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 8, pp. 1561–1576, Aug. 2011.

[29] Y. Feng, Y. Yang, and J. A. K. Suykens, "Robust gradient learning with applications," *IEEE Trans. Neural Netw. Learn. Syst.*, in press.

[30] J. Bolte, S. Sabach, and M. Teboulle, "Proximal alternating linearized minimization for nonconvex and nonsmooth problems," *Math. Program.*, vol. 146, nos. 1–2, pp. 459–494, 2014.

[31] T. G. Kolda and B. W. Bader, "Tensor decompositions and applications," *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, 2009.

[32] R. Tomioka, T. Suzuki, K. Hayashi, and H. Kashima, "Statistical performance of convex tensor decomposition," in *Proc. Adv. Neural Inf. Process. Syst.*, Granada, Spain, Dec. 2011, pp. 972–980.

[33] D. Kong, M. Zhang, and C. Ding, "Minimal shrinkage for noisy data recovery using Schatten-p norm objective," in *Machine Learning and Knowledge Discovery in Databases* (Lecture Notes in Computer Science), vol. 8189, H. Blockeel, K. Kersting, S. Nijssen, and F. Železný, Eds. Berlin, Germany: Springer-Verlag, 2013, pp. 177–193.

[34] E. J. Candès, X. Li, Y. Ma, and J. Wright, "Robust principal component analysis?" *J. ACM*, vol. 58, no. 3, Jan. 2011, Art. ID 11.

[35] H. Lu, K. N. Plataniotis, and A. N. Venetsanopoulos, "MPCA: Multilinear principal component analysis of tensor objects," *IEEE Trans. Neural Netw.*, vol. 19, no. 1, pp. 18–39, Jan. 2008.

[36] D. Goldfarb and S. Ma, "Convergence of fixed-point continuation algorithms for matrix rank minimization," *Found. Comput. Math.*, vol. 11, no. 2, pp. 183–210, 2011.

[37] R. T. Rockafellar and R. J.-B. Wets, *Variational Analysis*. Berlin, Germany: Springer-Verlag, 1998.

[38] J.-B. Hiriart-Urruty and H. Y. Le, "A variational approach of the rank function," *TOP*, vol. 21, no. 2, pp. 207–240, 2013.

[39] J.-F. Cai, E. J. Candès, and Z. Shen, "A singular value thresholding algorithm for matrix completion," *SIAM J. Optim.*, vol. 20, no. 4, pp. 1956–1982, 2010.

[40] S. Ma, D. Goldfarb, and L. Chen, "Fixed point and Bregman iterative methods for matrix rank minimization," *Math. Program.*, vol. 128, nos. 1–2, pp. 321–353, 2011.

[41] A. A. Goldstein, "Convex programming in Hilbert space," *Bull. Amer. Math. Soc.*, vol. 70, no. 5, pp. 709–710, 1964.

[42] L. Sorber, M. Van Barel, and L. De Lathauwer. *Tensorlab v2.0*. [Online]. Available: http://www.tensorlab.net/, accessed Jan. 2014.

[43] Y.-L. Chen, C.-T. Hsu, and H.-Y. M. Liao, "Simultaneous tensor decomposition and completion using factor priors," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 3, pp. 577–591, Mar. 2014.

[44] M. Burger, L. He, and C.-B. Schönlieb, "Cahn–Hilliard inpainting and a generalization for grayvalue images," *SIAM J. Imag. Sci.*, vol. 2, no. 4, pp. 1129–1167, 2009.

[45] Z. Lin, M. Chen, and Y. Ma. (Sep. 2010). "The augmented Lagrange multiplier method for exact recovery of corrupted low-rank matrices." [Online]. Available: http://arxiv.org/abs/1009.5055

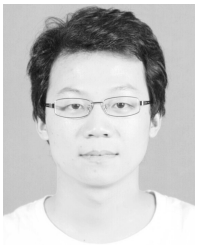
[46] J. Bochnak, M. Coste, and M.-F. Roy, *Real Algebraic Geometry*. New York, NY, USA: Springer-Verlag, 1998.

[47] M. Coste. (Oct. 2002). "An introduction to semialgebraic geometry," [Online]. Available: https://perso.univ-rennes1.fr/michel.coste/polyens/SAG.pdf



Yuning Yang was born in China in 1984. He received the B.S., M.S., and Ph.D. degrees in computational mathematics from Nankai University, Tianjin, China, in 2007, 2010, and 2013, respectively.

He has been a Post-Doctoral Researcher with Prof. J. A. K. Suykens with the Department of Electrical Engineering, Stadius Centre for Dynamical Systems, Signal Processing and Data Analytics, Katholieke Universiteit Leuven, Leuven, Belgium, since 2013. His current research interests include optimization, machine learning, and multilinear algebra.



Yunlong Feng was born in Anhui, China, in 1986. He received the bachelor's degree in mathematics from Nankai University, Tianjin, China, in 2007, and the Ph.D. degree in applied mathematics and computational mathematics from the City University of Hong Kong, Hong Kong, and the University of Science and Technology of China, Hefei, China, in 2012.

He was a Post-Doctoral Fellow with the Department of Mathematics, City University of Hong Kong, from 2012 to 2013. Since 2013, he has been a Post-Doctoral Researcher with Prof. J. A. K. Suykens with the Department of Electrical Engineering, Stadius Centre for Dynamical Systems, Signal Processing and Data Analytics, Katholieke Universiteit Leuven, Leuven, Belgium. His current research interests include statistical machine learning, learning theory, and nonparametric statistics.



Johan A. K. Suykens (F'15) was born in Willebroek, Belgium, in 1966. He received the master's degree in electromechanical engineering and the Ph.D. degree in applied sciences from the Katholieke Universiteit Leuven (KU Leuven), Leuven, Belgium, in 1989 and 1995, respectively.

He was a Visiting Post-Doctoral Researcher with the University of California at Berkeley, Berkeley, CA, USA, in 1996. He served as the Director and an Organizer of the NATO Advanced Study Institute on Learning Theory and Practice, Leuven, in 2002. He has been a Post-Doctoral Researcher with the Fund for Scientific Research FWO Flanders. He is currently a Professor (Hoogleraar) with KU Leuven. He has authored the books entitled *Artificial Neural Networks for Modeling and Control of Non-Linear Systems* (Kluwer Academic Publishers) and *Least Squares Support Vector Machines* (World Scientific) and co-authored a book entitled *Cellular Neural Networks, Multi-Scroll Chaos and Synchronization* (World Scientific), and edited the books entitled *Nonlinear Modeling: Advanced Black-Box Techniques* (Kluwer Academic Publishers), *Advances in Learning Theory: Methods, Models and Applications* (IOS Press), and *Regularization, Optimization, Kernels, and Support Vector Machines* (Chapman & Hall/CRC).

Prof. Suykens received the IEEE Signal Processing Society Best Paper Award in 1999 and several best paper awards at international conferences. He was a recipient of the International Neural Networks Society Young Investigator Award for his significant contributions in the field of neural networks in 2000. He also received the ERC Advanced Grant in 2011. In 1998, he organized an International Workshop on Nonlinear Modeling with Time-Series Prediction Competition. He served as an Associate Editor of the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS from 1997 to 1999 and 2004 to 2007, and the IEEE TRANSACTIONS ON NEURAL NETWORKS from 1998 to 2009. He served as the Program Co-Chair of the International Joint Conference on Neural Networks in 2004 and the International Symposium on Nonlinear Theory and its Applications in 2005, an Organizer of the International Symposium on Synchronization in Complex Networks in 2007, a Co-Organizer of the NIPS Workshop on Tensors, Kernels and Machine Learning in 2010.