

Parallel Algorithms for Constrained Tensor Factorization via Alternating Direction Method of Multipliers

Athanasios P. Liavas, *Member, IEEE*, and Nicholas D. Sidiropoulos, *Fellow, IEEE*

Abstract—Tensor factorization has proven useful in a wide range of applications, from sensor array processing to communications, speech and audio signal processing, and machine learning. With few recent exceptions, all tensor factorization algorithms were originally developed for centralized, in-memory computation on a single machine; and the few that break away from this mold do not easily incorporate practically important constraints, such as non-negativity. A new constrained tensor factorization framework is proposed in this paper, building upon the Alternating Direction Method of Multipliers (ADMoM). It is shown that this simplifies computations, bypassing the need to solve constrained optimization problems in each iteration; and it naturally leads to distributed algorithms suitable for parallel implementation. This opens the door for many emerging big data-enabled applications. The methodology is exemplified using non-negativity as a baseline constraint, but the proposed framework can incorporate many other types of constraints. Numerical experiments are encouraging, indicating that ADMoM-based non-negative tensor factorization (NTF) has high potential as an alternative to state-of-the-art approaches.

Index Terms—Tensor decomposition, PARAFAC model, parallel algorithms.

I. INTRODUCTION

TENSOR factorization¹ has proven useful in a wide range of signal processing applications, such as direction of arrival estimation [2], communication signal intelligence [3], and speech and audio signal separation [4], [5], as well as cross-disciplinary areas, such as community detection in social networks

[6], and chemical signal analysis [7]. More recently, there has been significant activity in applying tensor factorization theory and methods to problems in machine learning research—see [8].

There are two basic tensor factorization models: *parallel factor analysis* (PARAFAC) [9], [10] also known as *canonical decomposition* (CANDECOMP) [11], or CP (and CPD) for CANDECOMP-PARAFAC (Decomposition), or *canonical polyadic decomposition* (CPD, again); and the Tucker3 model [12]. Both are sum-of-outer-products models which historically served as cornerstones for further developments, e.g., block term decomposition [13], and upon which the vast majority of tensor applications have been built. In this paper, we will primarily focus on the CP model.

Whereas for low-enough² rank CP is already unique ‘on its own,’ any side information can (and should) be used to enhance identifiability and estimation performance in practice. Towards this end, we may exploit known properties of the sought latent factors, such as non-negativity, sparsity, monotonicity, or unimodality [14]. Whereas many of these properties can be handled with existing tensor factorization software, they generally complicate and slow down model fitting.

Unconstrained tensor factorization is already a hard non-convex (multi-linear) problem; even rank-one least-squares tensor approximation is NP-hard [15]. Many tensor factorization algorithms rely on alternating optimization, usually alternating least-squares (ALS), and imposing e.g., non-negativity and/or sparsity entails replacing linear least-squares conditional updates of the factor matrices with non-negative and/or sparse least-squares updates. In addition to ALS, many derivative-based methods have been developed that update all model parameters at once, see [16] and references therein, and [17], [18] for recent work in this direction.

With few recent exceptions, all tensor factorization algorithms were originally developed for centralized, in-memory computation on a single machine. This model of computation is inadequate for emerging big data-enabled applications, where the tensors to be analyzed cannot be loaded on a single machine, the data is more likely to reside in cloud storage, and cloud computing, or some other kind of high performance parallel architecture, must be used for the actual computation.

A carefully optimized Hadoop/MapReduce [19], [20] implementation of the basic ALS CP-decomposition algorithm was

Manuscript received August 30, 2014; revised January 21, 2015, May 05, 2015, and June 17, 2015; accepted June 23, 2015. Date of publication July 08, 2015; date of current version September 10, 2015. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Martin Haardt. The work of N. D. Sidiropoulos was supported in part by NSF IIS-1247632. Part of this work has been accepted for presentation at IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2015 [1].

A. P. Liavas is with the Department of Electronic and Computer Engineering, Technical University of Crete, Chania 73100, Greece (e-mail: liavas@telecom.tuc.gr).

N. D. Sidiropoulos is with the Department of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN 55455 USA (e-mail: nikos@ece.umn.edu).

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/TSP.2015.2454476

¹In the literature, the terms *factorization* and *decomposition* are often used interchangeably, even though the latter alludes to exact decomposition, whereas the former may include a residual term.

²E.g., relative to the sum of Kruskal-ranks of the latent factor matrices. Looser bounds can be guaranteed almost-surely.

developed in [21], which reported 100-fold scaling improvements relative to the prior art. The gist of [21] is to avoid the explicit computation of ‘blown-up’ intermediate matrix products in the ALS algorithm, particularly for sparse tensors, and parallelization is achieved by splitting the computation of outer products. On the other hand, [21] is not designed for high performance computing (e.g., mesh) architectures, and it does not incorporate constraints on the factor matrices.

A random sampling approach was later proposed in [6], motivated by recent progress in randomized algorithms for matrix algebra. The idea of [6] is to create and analyze multiple randomly sub-sampled parts of the tensor, then combine the results using a common piece of data to anchor the constituent decompositions. The downside of [6] is that it only works for sparse tensors, and it offers no identifiability guarantees—although it usually works well for sparse tensors.

A different approach based on generalized random sampling was recently proposed in [22], [23]. The idea is to create multiple randomly compressed mixtures (instead of sub-sampled parts) of the original tensor, analyze them all in parallel, and then combine the results. The main advantages of [22], [23] over [6] are that i) identifiability can be guaranteed, ii) no sparsity is needed, and iii) there are theoretical scalability guarantees.

Distributed CP decomposition based on the ALS algorithm has been considered in [24], and more recently in [25], which exploit the inherent parallelism in the matrix version of the linear least-squares subproblems to split the computation in different ways, assuming an essentially ‘flat’ architecture for the computing nodes. Regular (e.g., mesh) architectures and constraints on the latent factors are not considered in [24], [25].

In this paper, we develop algorithms for constrained tensor factorization based on Alternating Direction Method of Multipliers (ADMoM). ADMoM has recently attracted renewed interest [26], primarily for solving certain types of convex optimization problems in a distributed fashion. However, it can also be used to tackle non-convex problems, such as non-negative matrix factorization [26], albeit its convergence properties are far less understood in this case. We focus on non-negative CP decompositions as a working problem, due to the importance of the CP model and non-negativity constraints; but our approach can be generalized to many other types of constraints on the latent factors, as well as other tensor factorizations, such as Tucker3, and tensor completion.

The advantages of our approach are as follows. First, during each ADMoM iteration, we avoid the solution of constrained optimization problems, resulting in considerably smaller computational complexity per iteration compared to constrained least-squares based algorithms, such as alternating non-negative least-squares (NALS). Second, our approach leads naturally to distributed algorithms suitable for parallel implementation on regular high-performance computing (e.g., mesh) architectures. Finally, our approach can easily incorporate many other types of constraints on the latent factors, such as sparsity.

Numerical experiments are encouraging, indicating that ADMoM-based NTF has significant potential as an alternative to state-of-the-art approaches.

The rest of the manuscript is structured as follows. In Section II, we present the NTF problem and in Section III we

present the general ADMoM framework. In Section IV, we develop ADMoM for NTF, while in Section V we develop distributed ADMoM for large NTF. In Section VI, we test the behavior of the developed schemes with numerical experiments. Finally, in Section VII, we conclude the paper.

Notation

Vectors, matrices, and tensors are denoted by small, capital, and underlined capital bold letters, respectively; for example, \mathbf{x} , \mathbf{X} , and $\underline{\mathbf{X}}$. $\mathbb{R}_+^{I \times J \times K}$ denotes the set of $(I \times J \times K)$ real non-negative tensors, while $\mathbb{R}_+^{I \times J}$ denotes the set of $(I \times J)$ real non-negative matrices. $\|\cdot\|_F$ denotes the Frobenius norm of the tensor or matrix argument, \mathbf{A}^\dagger denotes the Moore-Penrose pseudoinverse of matrix \mathbf{A} , and $(\mathbf{A})_+$ denotes the projection of matrix \mathbf{A} onto the set of element-wise non-negative matrices. The outer product of three vectors $\mathbf{a} \in \mathbb{R}^{I \times 1}$, $\mathbf{b} \in \mathbb{R}^{J \times 1}$, and $\mathbf{c} \in \mathbb{R}^{K \times 1}$ is the rank-one tensor $\mathbf{a} \circ \mathbf{b} \circ \mathbf{c} \in \mathbb{R}^{I \times J \times K}$ with elements $(\mathbf{a} \circ \mathbf{b} \circ \mathbf{c})(i, j, k) = \mathbf{a}(i)\mathbf{b}(j)\mathbf{c}(k)$. For matrices \mathbf{A} and \mathbf{B} , with compatible dimensions, $\mathbf{A} \odot \mathbf{B}$ denotes the Khatri-Rao (columnwise Kronecker) product, $\mathbf{A} \circledast \mathbf{B}$ denotes the Hadamard (element-wise) product, and $\mathbf{A} * \mathbf{B}$ denotes the matrix inner product, that is $\mathbf{A} * \mathbf{B} := \text{trace}(\mathbf{A}^T \mathbf{B}) = \sum_{i,j} \mathbf{A}_{i,j} \mathbf{B}_{i,j}$.

II. NON-NEGATIVE TENSOR FACTORIZATION

Let tensor $\underline{\mathbf{X}}^\circ \in \mathbb{R}_+^{I \times J \times K}$ admit a non-negative³ CP decomposition of order F

$$\underline{\mathbf{X}}^\circ = [\mathbf{A}^\circ, \mathbf{B}^\circ, \mathbf{C}^\circ] = \sum_{f=1}^F \mathbf{a}_f^\circ \circ \mathbf{b}_f^\circ \circ \mathbf{c}_f^\circ,$$

where $\mathbf{A}^\circ = [\mathbf{a}_1^\circ \cdots \mathbf{a}_F^\circ] \in \mathbb{R}_+^{I \times F}$, $\mathbf{B}^\circ = [\mathbf{b}_1^\circ \cdots \mathbf{b}_F^\circ] \in \mathbb{R}_+^{J \times F}$, and $\mathbf{C}^\circ = [\mathbf{c}_1^\circ \cdots \mathbf{c}_F^\circ] \in \mathbb{R}_+^{K \times F}$. We observe a noisy version of $\underline{\mathbf{X}}^\circ$ expressed as

$$\underline{\mathbf{X}} = \underline{\mathbf{X}}^\circ + \underline{\mathbf{E}}.$$

In order to estimate \mathbf{A}° , \mathbf{B}° , and \mathbf{C}° , we compute matrices $\mathbf{A} \in \mathbb{R}_+^{I \times F}$, $\mathbf{B} \in \mathbb{R}_+^{J \times F}$, and $\mathbf{C} \in \mathbb{R}_+^{K \times F}$ that solve the optimization problem

$$\begin{aligned} \min_{\mathbf{A}, \mathbf{B}, \mathbf{C}} \quad & f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) \\ \text{subject to} \quad & \mathbf{A} \geq \mathbf{0}, \mathbf{B} \geq \mathbf{0}, \mathbf{C} \geq \mathbf{0}, \end{aligned} \quad (1)$$

where f is a function measuring the quality of the factorization, $\mathbf{0}$ is the zero matrix of appropriate dimensions, and the inequalities are element-wise. A common choice for $f_{\underline{\mathbf{X}}}$, motivated via maximum likelihood estimation for $\underline{\mathbf{E}}$ with Gaussian independent and identically distributed (i.i.d.) elements, is

$$f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) = \frac{1}{2} \|\underline{\mathbf{X}} - [\mathbf{A}, \mathbf{B}, \mathbf{C}]\|_F^2. \quad (2)$$

³Note that, due to the non-negativity constraints on the latent factors, F can be higher than the rank of $\underline{\mathbf{X}}^\circ$.

Let $\underline{\mathbf{W}} = [\mathbf{A}, \mathbf{B}, \mathbf{C}]$ and $\mathbf{W}^{(1)}$, $\mathbf{W}^{(2)}$, and $\mathbf{W}^{(3)}$ be the matrix unfoldings of $\underline{\mathbf{W}}$, with respect to the first, second, and third dimension, respectively. Then,

$$\begin{aligned}\mathbf{W}^{(1)} &= \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T, \\ \mathbf{W}^{(2)} &= \mathbf{B}(\mathbf{C} \odot \mathbf{A})^T, \\ \mathbf{W}^{(3)} &= \mathbf{C}(\mathbf{B} \odot \mathbf{A})^T,\end{aligned}\quad (3)$$

and $f_{\underline{\mathbf{X}}}$ can be equivalently expressed as

$$\begin{aligned}f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) &= \frac{1}{2} \left\| \mathbf{X}^{(1)} - \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T \right\|_F^2 \\ &= \frac{1}{2} \left\| \mathbf{X}^{(2)} - \mathbf{B}(\mathbf{C} \odot \mathbf{A})^T \right\|_F^2 \\ &= \frac{1}{2} \left\| \mathbf{X}^{(3)} - \mathbf{C}(\mathbf{B} \odot \mathbf{A})^T \right\|_F^2.\end{aligned}\quad (4)$$

These expressions are the basis for ALS-type CP optimization, because they enable simple linear least-squares updating of one matrix given the other two. Using NALS for each update step is a popular approach for the solution of (1), but non-negativity brings a significant computational burden relative to plain ALS and also complicates the development of parallel algorithms for NTF. It is worth noting that (4) will also prove useful during the development of the ADMoM-based NTF algorithm.

III. ADMoM

ADMoM is a technique for the solution of optimization problems of the form [26]

$$\begin{aligned}\min_{\mathbf{x}, \mathbf{z}} \quad & f(\mathbf{x}) + g(\mathbf{z}) \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} = \mathbf{c},\end{aligned}\quad (5)$$

where $\mathbf{x} \in \mathbb{R}^{n_1}$, $\mathbf{z} \in \mathbb{R}^{n_2}$, $\mathbf{A} \in \mathbb{R}^{m \times n_1}$, $\mathbf{B} \in \mathbb{R}^{m \times n_2}$, $\mathbf{c} \in \mathbb{R}^m$, $f: \mathbb{R}^{n_1} \rightarrow \mathbb{R}$, and $g: \mathbb{R}^{n_2} \rightarrow \mathbb{R}$.

The augmented Lagrangian for problem (5) is

$$\begin{aligned}L_{\rho}(\mathbf{x}, \mathbf{z}, \mathbf{y}) &= f(\mathbf{x}) + g(\mathbf{z}) + \mathbf{y}^T(\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c}) \\ &\quad + \frac{\rho}{2} \|\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z} - \mathbf{c}\|_2^2,\end{aligned}\quad (6)$$

where $\rho > 0$ is a penalty parameter. Assuming that at time instant k we have computed \mathbf{z}^k and \mathbf{y}^k , which comprise the state of the algorithm, the $(k+1)$ -st iteration of ADMoM is⁴

$$\begin{aligned}\mathbf{x}^{k+1} &= \arg \min_{\mathbf{x}} \left(f(\mathbf{x}) + \mathbf{y}^{kT} \mathbf{A}\mathbf{x} + \frac{\rho}{2} \|\mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{z}^k - \mathbf{c}\|_2^2 \right) \\ \mathbf{z}^{k+1} &= \arg \min_{\mathbf{z}} \left(g(\mathbf{z}) + \mathbf{y}^{kT} \mathbf{B}\mathbf{z} + \frac{\rho}{2} \|\mathbf{A}\mathbf{x}^{k+1} + \mathbf{B}\mathbf{z} - \mathbf{c}\|_2^2 \right) \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + \rho(\mathbf{A}\mathbf{x}^{k+1} + \mathbf{B}\mathbf{z}^{k+1} - \mathbf{c}).\end{aligned}$$

It can be shown that, under certain conditions (among them convexity of f and g), ADMoM converges in a certain sense (see [26] for an excellent review of ADMoM, including some convergence analysis results). However, ADMoM can be used even when problem (5) is *non-convex*. In this case, we use ADMoM with the goal of reaching a good local minimum [26]. Note

⁴Note that \mathbf{y}^{kT} is shorthand notation for $(\mathbf{y}^k)^T$.

that this is all we can realistically hope for anyway, irrespective of approach or algorithm used, since tensor factorization is NP-hard [15].

A. ADMoM for Set-Constrained Optimization

Let us consider the set-constrained optimization problem

$$\begin{aligned}\min_{\mathbf{x}} \quad & f(\mathbf{x}) \\ \text{subject to} \quad & \mathbf{x} \in \mathcal{X},\end{aligned}$$

where $\mathbf{x} \in \mathbb{R}^n$ and $\mathcal{X} \subseteq \mathbb{R}^n$ is a closed convex set. At first sight, this problem does not seem suitable for ADMoM. However, if we introduce variable $\mathbf{z} \in \mathbb{R}^n$, we can write the equivalent problem

$$\begin{aligned}\min_{\mathbf{x}, \mathbf{z}} \quad & f(\mathbf{x}) + g(\mathbf{z}) \\ \text{subject to} \quad & \mathbf{x} - \mathbf{z} = \mathbf{0},\end{aligned}\quad (7)$$

where g is the indicator function of set \mathcal{X} , that is,

$$g(\mathbf{z}) := \begin{cases} 0, & \mathbf{z} \in \mathcal{X}, \\ \infty, & \mathbf{z} \notin \mathcal{X}. \end{cases}$$

Then it becomes clear that (7) can be solved via ADMoM. Assuming that at time instant k we have computed \mathbf{z}^k and \mathbf{y}^k , the $(k+1)$ -st iteration of ADMoM is [26]

$$\begin{aligned}\mathbf{x}^{k+1} &= \arg \min_{\mathbf{x}} \left(f(\mathbf{x}) + \mathbf{y}^{kT} \mathbf{x} + \frac{\rho}{2} \|\mathbf{x} - \mathbf{z}^k\|_2^2 \right) \\ \mathbf{z}^{k+1} &= \arg \min_{\mathbf{z}} \left(g(\mathbf{z}) - \mathbf{y}^{kT} \mathbf{z} + \frac{\rho}{2} \|\mathbf{x}^{k+1} - \mathbf{z}\|_2^2 \right) \\ &= \Pi_{\mathcal{X}} \left(\mathbf{x}^{k+1} + \frac{1}{\rho} \mathbf{y}^k \right) \\ \mathbf{y}^{k+1} &= \mathbf{y}^k + \rho(\mathbf{x}^{k+1} - \mathbf{z}^{k+1}),\end{aligned}$$

where $\Pi_{\mathcal{X}}$ denotes projection (in the Euclidean norm) onto \mathcal{X} .

IV. ADMoM FOR NTF

In this section, we adopt the approach of Section III-A and develop an ADMoM-based NTF algorithm. At first, we must put the NTF problem (1) into ADMoM form. Towards this end, we introduce auxiliary variables $\tilde{\mathbf{A}} \in \mathbb{R}^{I \times F}$, $\tilde{\mathbf{B}} \in \mathbb{R}^{J \times F}$, and $\tilde{\mathbf{C}} \in \mathbb{R}^{K \times F}$ and consider the equivalent optimization problem

$$\begin{aligned}\min_{\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{B}, \tilde{\mathbf{B}}, \mathbf{C}, \tilde{\mathbf{C}}} \quad & f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) + g(\tilde{\mathbf{A}}) + g(\tilde{\mathbf{B}}) + g(\tilde{\mathbf{C}}) \\ \text{subject to} \quad & \mathbf{A} - \tilde{\mathbf{A}} = \mathbf{0}, \mathbf{B} - \tilde{\mathbf{B}} = \mathbf{0}, \mathbf{C} - \tilde{\mathbf{C}} = \mathbf{0},\end{aligned}\quad (8)$$

where, for any matrix argument \mathbf{M} ,

$$g(\mathbf{M}) := \begin{cases} 0, & \text{if } \mathbf{M} \geq \mathbf{0}, \\ \infty, & \text{otherwise.} \end{cases}\quad (9)$$

We introduce the dual variables $\mathbf{Y}_{\mathbf{A}} \in \mathbb{R}^{I \times F}$, $\mathbf{Y}_{\mathbf{B}} \in \mathbb{R}^{J \times F}$, and $\mathbf{Y}_{\mathbf{C}} \in \mathbb{R}^{K \times F}$, and the vector of penalty terms

$\rho := [\rho_A \ \rho_B \ \rho_C]^T$. The augmented Lagrangian is given in (10), at the bottom of the page.

The ADMoM for problem (8) is as follows:

$$\begin{aligned}
 (\mathbf{A}^{k+1}, \mathbf{B}^{k+1}, \mathbf{C}^{k+1}) &= \arg \min_{\mathbf{A}, \mathbf{B}, \mathbf{C}} \left(f_{\mathbf{X}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) \right. \\
 &\quad + \mathbf{Y}_A^k * \mathbf{A} + \frac{\rho_A}{2} \|\mathbf{A} - \tilde{\mathbf{A}}^k\|_F^2 \\
 &\quad + \mathbf{Y}_B^k * \mathbf{B} + \frac{\rho_B}{2} \|\mathbf{B} - \tilde{\mathbf{B}}^k\|_F^2 \\
 &\quad \left. + \mathbf{Y}_C^k * \mathbf{C} + \frac{\rho_C}{2} \|\mathbf{C} - \tilde{\mathbf{C}}^k\|_F^2 \right) \\
 \tilde{\mathbf{A}}^{k+1} &= \left(\mathbf{A}^{k+1} + \frac{1}{\rho_A} \mathbf{Y}_A^k \right)_+ \\
 \tilde{\mathbf{B}}^{k+1} &= \left(\mathbf{B}^{k+1} + \frac{1}{\rho_B} \mathbf{Y}_B^k \right)_+ \\
 \tilde{\mathbf{C}}^{k+1} &= \left(\mathbf{C}^{k+1} + \frac{1}{\rho_C} \mathbf{Y}_C^k \right)_+ \\
 \mathbf{Y}_A^{k+1} &= \mathbf{Y}_A^k + \rho_A (\mathbf{A}^{k+1} - \tilde{\mathbf{A}}^{k+1}) \\
 \mathbf{Y}_B^{k+1} &= \mathbf{Y}_B^k + \rho_B (\mathbf{B}^{k+1} - \tilde{\mathbf{B}}^{k+1}) \\
 \mathbf{Y}_C^{k+1} &= \mathbf{Y}_C^k + \rho_C (\mathbf{C}^{k+1} - \tilde{\mathbf{C}}^{k+1}). \quad (11)
 \end{aligned}$$

The minimization problem in the first line of (11) is non-convex. Using the equivalent expressions for $f_{\mathbf{X}}$ in (4), we propose the alternating optimization scheme of (12) at the bottom of the page. The updates of (12) can be executed either for a predetermined number of iterations, or until convergence.⁵ We observe

⁵In our implementations, we execute these updates once per ADMoM iteration.

that, during each ADMoM iteration, we avoid the solution of constrained optimization problems. This seems favorable, especially in the cases where the size of the problem is (very) large.

We note that \mathbf{A}^k , \mathbf{B}^k , and \mathbf{C}^k are *not* necessarily non-negative. They become non-negative (or, at least, their negative elements become very small) upon convergence. On the other hand, $\tilde{\mathbf{A}}^k$, $\tilde{\mathbf{B}}^k$, and $\tilde{\mathbf{C}}^k$ are by construction non-negative.

A. Computational Complexity per Iteration

Each ADMoM iteration consists of simple matrix operations. Thus, rough estimates of its computational complexity can be easily derived (of course, accurate estimates can be derived after fixing the algorithms that implement the matrix operations).

A rough estimate for the computational complexity of the update of \mathbf{A}^k (see the first update in (12)) can be derived as follows:

- 1) $O(nF + IF)$ for the computation of the term $\mathbf{X}^{(1)}(\mathbf{C}^k \odot \mathbf{B}^k) + \rho_A \tilde{\mathbf{A}}^k - \mathbf{Y}_A^k$, where n is the number of nonzero elements of tensor \mathbf{X} (also of matrix $\mathbf{X}^{(1)}$). Note that $n = IJK$ for dense tensors, but for sparse tensors $n \ll IJK$. This is because the product $\mathbf{X}^{(1)}(\mathbf{C}^k \odot \mathbf{B}^k)$ can be computed with $3nF$ flops, by exploiting sparsity and the structure of the Khatri-Rao product [27]–[29]. With $5nF$ flops, it is possible to parallelize this computation [21]. Efficient (in terms of favorable memory access pattern) in-place computation of all three products needed for the update of \mathbf{A}^k , \mathbf{B}^k , \mathbf{C}^k from a single copy of \mathbf{X} has been recently considered in [30], which also features potential flop gains as a side-benefit.
- 2) $O((K + J)F^2)$ for the computation of the term $(\mathbf{C}^k \odot \mathbf{B}^k)^T(\mathbf{C}^k \odot \mathbf{B}^k) + \rho_A \mathbf{I}_F$, and $O(F^3)$ for its Cholesky

$$\begin{aligned}
 L_{\rho}(\mathbf{A}, \mathbf{B}, \mathbf{C}, \tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}, \mathbf{Y}_A, \mathbf{Y}_B, \mathbf{Y}_C) &= f_{\mathbf{X}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) + g(\tilde{\mathbf{A}}) + g(\tilde{\mathbf{B}}) + g(\tilde{\mathbf{C}}) \\
 &\quad + \mathbf{Y}_A * (\mathbf{A} - \tilde{\mathbf{A}}) + \frac{\rho_A}{2} \|\mathbf{A} - \tilde{\mathbf{A}}\|_F^2 \\
 &\quad + \mathbf{Y}_B * (\mathbf{B} - \tilde{\mathbf{B}}) + \frac{\rho_B}{2} \|\mathbf{B} - \tilde{\mathbf{B}}\|_F^2 \\
 &\quad + \mathbf{Y}_C * (\mathbf{C} - \tilde{\mathbf{C}}) + \frac{\rho_C}{2} \|\mathbf{C} - \tilde{\mathbf{C}}\|_F^2. \quad (10)
 \end{aligned}$$

$$\begin{aligned}
 \mathbf{A}^{k+1} &= \arg \min_{\mathbf{A}} \left(\frac{1}{2} \|\mathbf{X}^{(1)} - \mathbf{A}(\mathbf{C}^k \odot \mathbf{B}^k)^T\|_F^2 + \mathbf{Y}_A^k * \mathbf{A} + \frac{\rho_A}{2} \|\mathbf{A} - \tilde{\mathbf{A}}^k\|_F^2 \right) \\
 &= \left(\mathbf{X}^{(1)}(\mathbf{C}^k \odot \mathbf{B}^k) + \rho_A \tilde{\mathbf{A}}^k - \mathbf{Y}_A^k \right) \left((\mathbf{C}^k \odot \mathbf{B}^k)^T(\mathbf{C}^k \odot \mathbf{B}^k) + \rho_A \mathbf{I}_F \right)^{-1}, \\
 \mathbf{B}^{k+1} &= \arg \min_{\mathbf{B}} \left(\frac{1}{2} \|\mathbf{X}^{(2)} - \mathbf{B}(\mathbf{C}^k \odot \mathbf{A}^{k+1})^T\|_F^2 + \mathbf{Y}_B^k * \mathbf{B} + \frac{\rho_B}{2} \|\mathbf{B} - \tilde{\mathbf{B}}^k\|_F^2 \right) \\
 &= \left(\mathbf{X}^{(2)}(\mathbf{C}^k \odot \mathbf{A}^{k+1}) + \rho_B \tilde{\mathbf{B}}^k - \mathbf{Y}_B^k \right) \left((\mathbf{C}^k \odot \mathbf{A}^{k+1})^T(\mathbf{C}^k \odot \mathbf{A}^{k+1}) + \rho_B \mathbf{I}_F \right)^{-1}, \\
 \mathbf{C}^{k+1} &= \arg \min_{\mathbf{C}} \left(\frac{1}{2} \|\mathbf{X}^{(3)} - \mathbf{C}(\mathbf{B}^{k+1} \odot \mathbf{A}^{k+1})^T\|_F^2 + \mathbf{Y}_C^k * \mathbf{C} + \frac{\rho_C}{2} \|\mathbf{C} - \tilde{\mathbf{C}}^k\|_F^2 \right) \\
 &= \left(\mathbf{X}^{(3)}(\mathbf{B}^{k+1} \odot \mathbf{A}^{k+1}) + \rho_C \tilde{\mathbf{C}}^k - \mathbf{Y}_C^k \right) \left((\mathbf{B}^{k+1} \odot \mathbf{A}^{k+1})^T(\mathbf{B}^{k+1} \odot \mathbf{A}^{k+1}) + \rho_C \mathbf{I}_F \right)^{-1}. \quad (12)
 \end{aligned}$$

decomposition. This is because $(\mathbf{C}^k \odot \mathbf{B}^k)^T (\mathbf{C}^k \odot \mathbf{B}^k) = ((\mathbf{C}^k)^T \mathbf{C}^k) \circ ((\mathbf{B}^k)^T \mathbf{B}^k)$.

- 3) $O(F^2 I)$ for the computation of the system solution that gives the updated value \mathbf{A}^{k+1} .

Analogous estimates can be derived for the updates of \mathbf{B}^k and \mathbf{C}^k . Finally, the updates of the auxiliary and dual variables require, in total, $O((I + J + K)F)$ arithmetic operations.

B. Convergence

Let $\mathbf{Z} := (\mathbf{A}, \mathbf{B}, \mathbf{C}, \tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}, \mathbf{Y}_\mathbf{A}, \mathbf{Y}_\mathbf{B}, \mathbf{Y}_\mathbf{C})$. It can be proven that \mathbf{Z} is a Karush-Kuhn-Tucker (KKT) point for the NTF problem (8) if

$$\begin{aligned} (\mathbf{X}^{(1)} - \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T)(\mathbf{C} \odot \mathbf{B}) - \mathbf{Y}_\mathbf{A} &= \mathbf{0} \\ (\mathbf{X}^{(2)} - \mathbf{B}(\mathbf{C} \odot \mathbf{A})^T)(\mathbf{C} \odot \mathbf{A}) - \mathbf{Y}_\mathbf{B} &= \mathbf{0} \\ (\mathbf{X}^{(3)} - \mathbf{C}(\mathbf{B} \odot \mathbf{A})^T)(\mathbf{B} \odot \mathbf{A}) - \mathbf{Y}_\mathbf{C} &= \mathbf{0} \\ \mathbf{A} - \tilde{\mathbf{A}} = \mathbf{0}, \mathbf{B} - \tilde{\mathbf{B}} = \mathbf{0}, \mathbf{C} - \tilde{\mathbf{C}} &= \mathbf{0} \\ \mathbf{Y}_\mathbf{A} \leq \mathbf{0}, \mathbf{Y}_\mathbf{B} \leq \mathbf{0}, \mathbf{Y}_\mathbf{C} \leq \mathbf{0} \\ \mathbf{Y}_\mathbf{A} \circ \tilde{\mathbf{A}} = \mathbf{0}, \mathbf{Y}_\mathbf{B} \circ \tilde{\mathbf{B}} = \mathbf{0}, \mathbf{Y}_\mathbf{C} \circ \tilde{\mathbf{C}} &= \mathbf{0}. \end{aligned} \quad (13)$$

Proposition 1: Let $\{\mathbf{Z}^k\}$ be a sequence generated by ADMoM for NTF that satisfies condition

$$\lim_{k \rightarrow \infty} (\mathbf{Z}^{k+1} - \mathbf{Z}^k) = \mathbf{0}. \quad (14)$$

Then, any accumulation point of $\{\mathbf{Z}^k\}$ is a KKT point of problem (8). Consequently, any accumulation point of $\{\mathbf{A}^k, \mathbf{B}^k, \mathbf{C}^k\}$ is a KKT point of problem (1).

Proof: The proof follows closely the steps of the proof of Proposition 2.1 of [31] and is omitted. \square

Proposition 1 implies that, whenever $\{\mathbf{Z}^k\}$ converges, it converges to a KKT point. We will further discuss ADMoM convergence from a practical point of view in the section with the numerical experiments.

C. Stopping Criteria

The primal residual for variable \mathbf{A}^k is defined as

$$\mathbf{P}_\mathbf{A}^k := \mathbf{A}^k - \tilde{\mathbf{A}}^k, \quad (15)$$

while quantity

$$\mathbf{D}_\mathbf{A}^k := \rho_\mathbf{A}(\tilde{\mathbf{A}}^k - \tilde{\mathbf{A}}^{k-1}) \quad (16)$$

can be viewed as a dual feasibility residual (see [26, Section 3.3]). We analogously define $\mathbf{P}_\mathbf{B}^k$, $\mathbf{D}_\mathbf{B}^k$, $\mathbf{P}_\mathbf{C}^k$, and $\mathbf{D}_\mathbf{C}^k$.

We stop the algorithm if all primal and dual residuals are sufficiently small. More specifically, we introduce small positive constants ϵ^{abs} and ϵ^{rel} and consider $\mathbf{P}_\mathbf{A}^k$ and $\mathbf{D}_\mathbf{A}^k$ small if

$$\|\mathbf{P}_\mathbf{A}^k\|_F \leq \sqrt{IF} \epsilon^{\text{abs}} + \epsilon^{\text{rel}} \max \left\{ \|\mathbf{A}^k\|_F, \|\tilde{\mathbf{A}}^k\|_F \right\}, \quad (17)$$

$$\|\mathbf{D}_\mathbf{A}^k\|_F \leq \sqrt{IF} \epsilon^{\text{abs}} + \epsilon^{\text{rel}} \|\mathbf{Y}_\mathbf{A}\|_F. \quad (18)$$

⁶See report [32] for a detailed proof.

Analogous conditions apply for the other residuals. Reasonable values for ϵ^{rel} are $\epsilon^{\text{rel}} \lesssim 10^{-3}$, while the value of ϵ^{abs} depends on the scale of the values of the latent factors.

We note that stopping criteria (17) and (18) involve quantities of the size of the latent factors which, in most cases, is small compared to the size of the tensor. Thus, their computation, even during every ADMoM iteration, is not computationally demanding.

D. Varying Penalty Parameters

We have found very useful in practice to vary the values of each one of the penalty parameters, $\rho_\mathbf{A}$, $\rho_\mathbf{B}$, and $\rho_\mathbf{C}$, depending on the size of the corresponding primal and dual residuals (see [26, Section 3.4]). More specifically, the penalty parameters $\rho_\mathbf{M}^k$, for $\mathbf{M} = \mathbf{A}, \mathbf{B}, \mathbf{C}$, are updated as follows:

$$\rho_\mathbf{M}^{k+1} = \begin{cases} \tau^{\text{incr}} \rho_\mathbf{M}^k, & \text{if } \|\mathbf{P}_\mathbf{M}^k\|_F > \mu \|\mathbf{D}_\mathbf{M}^k\|_F, \\ \rho_\mathbf{M}^k / \tau^{\text{decr}}, & \text{if } \|\mathbf{D}_\mathbf{M}^k\|_F > \mu \|\mathbf{P}_\mathbf{M}^k\|_F, \\ \rho_\mathbf{M}^k, & \text{otherwise,} \end{cases} \quad (19)$$

where $\mu > 1$, $\tau^{\text{incr}} > 1$, and $\tau^{\text{decr}} > 1$ are the adaptation parameters. Large values of $\rho_\mathbf{M}$ place large penalty on violations of primal feasibility, leading to small primal residuals, while small values of $\rho_\mathbf{M}$ tend to reduce the dual residuals.

E. ADMoM for Tensor Factorization With Structural Constraints

ADMoM can easily handle certain structural constraints on the latent factors [26], [33]. For example, if we want to solve an NTF problem with the added constraint that the number of nonzero elements of \mathbf{A} is lower than or equal to a given number $c_\mathbf{A}$, then we can adopt an approach similar to that followed in Section IV with the only difference being that, instead of using $g(\tilde{\mathbf{A}})$ defined in (9), we use $g_{c_\mathbf{A}}(\tilde{\mathbf{A}})$ where, for any matrix argument \mathbf{M} ,

$$g_c(\mathbf{M}) := \begin{cases} 0, & \text{if } \mathbf{M} \geq \mathbf{0} \text{ and } \|\mathbf{M}\|_0 \leq c, \\ \infty, & \text{otherwise.} \end{cases} \quad (20)$$

The only difference between the ADMoM for this case and the one presented in (11) and (12) is in the update of $\tilde{\mathbf{A}}^k$. More specifically, instead of using projection onto the set of non-negative matrices, we must use projection onto the set of non-negative matrices with at most $c_\mathbf{A}$ nonzero elements, which can be easily computed through sorting of the elements of $\tilde{\mathbf{A}}^k$. Using analogous arguments, we can incorporate into our ADMoM framework box or other set constraints on the latent factors. The development of the corresponding ADMoM is almost trivial if projection onto the constraint set is easy.

Thorough study of ADMoM-based algorithms for tensor factorization and/or completion with more complicated structural constraints is a topic of future research.

V. DISTRIBUTED ADMOM FOR LARGE NTF

In this section, we assume that all dimensions of tensor $\underline{\mathbf{X}}$ are large and derive an ADMoM-based NTF that is suitable for parallel implementation. Of course, our framework can handle the cases where only one or two of the dimensions of $\underline{\mathbf{X}}$ are large.

A. Matrix Unfoldings in Terms of Partitioned Matrix Factors

Let $\underline{\mathbf{W}} = [\mathbf{A}, \mathbf{B}, \mathbf{C}]$, and \mathbf{A} , \mathbf{B} , and \mathbf{C} be partitioned as

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_{N_A} \end{bmatrix}, \mathbf{B} = \begin{bmatrix} \mathbf{B}_1 \\ \vdots \\ \mathbf{B}_{N_B} \end{bmatrix}, \mathbf{C} = \begin{bmatrix} \mathbf{C}_1 \\ \vdots \\ \mathbf{C}_{N_C} \end{bmatrix}, \quad (21)$$

with $\mathbf{A}_{n_A} \in \mathbb{R}^{I_{n_A} \times F}$, for $n_A = 1, \dots, N_A$, $\sum_{n_A=1}^{N_A} I_{n_A} = I$, $\mathbf{B}_{n_B} \in \mathbb{R}^{J_{n_B} \times F}$, for $n_B = 1, \dots, N_B$, $\sum_{n_B=1}^{N_B} J_{n_B} = J$, and $\mathbf{C}_{n_C} \in \mathbb{R}^{K_{n_C} \times F}$, for $n_C = 1, \dots, N_C$, $\sum_{n_C=1}^{N_C} K_{n_C} = K$.

We first derive partitionings of the matrix unfoldings of $\underline{\mathbf{W}}$ in terms of (the blocks of) matrices \mathbf{A} , \mathbf{B} , and \mathbf{C} . Towards this end, we write $\mathbf{W}^{(1)}$ as

$$\begin{aligned} \mathbf{W}^{(1)} &= \mathbf{A}(\mathbf{C} \odot \mathbf{B})^T \\ &= \begin{bmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_{N_A} \end{bmatrix} \left(\begin{bmatrix} \mathbf{C}_1 \\ \vdots \\ \mathbf{C}_{N_C} \end{bmatrix} \odot \mathbf{B} \right)^T \\ &= \begin{bmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_{N_A} \end{bmatrix} \left(\begin{bmatrix} \mathbf{C}_1 \odot \mathbf{B} \\ \vdots \\ \mathbf{C}_{N_C} \odot \mathbf{B} \end{bmatrix} \right)^T \\ &= \begin{bmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_{N_A} \end{bmatrix} \left([(\mathbf{C}_1 \odot \mathbf{B})^T \ \cdots \ (\mathbf{C}_{N_C} \odot \mathbf{B})^T] \right). \end{aligned}$$

Thus, $\mathbf{W}^{(1)}$ can be partitioned as

$$\mathbf{W}^{(1)} = \begin{bmatrix} \mathbf{W}_{1,1}^{(1)} & \cdots & \mathbf{W}_{1,N_C}^{(1)} \\ \vdots & \ddots & \vdots \\ \mathbf{W}_{N_A,1}^{(1)} & \cdots & \mathbf{W}_{N_A,N_C}^{(1)} \end{bmatrix},$$

where the (n_A, n_C) -th block of $\mathbf{W}^{(1)}$ is equal to the $I_{n_A} \times (JK_{n_C})$ matrix $\mathbf{W}_{n_A,n_C}^{(1)} = \mathbf{A}_{n_A}(\mathbf{C}_{n_C} \odot \mathbf{B})^T$, for $n_A = 1, \dots, N_A$ and $n_C = 1, \dots, N_C$.

Similarly, it can be shown that $\mathbf{W}^{(2)}$ can be partitioned into blocks $\mathbf{W}_{n_B,n_C}^{(2)} = \mathbf{B}_{n_B}(\mathbf{C}_{n_C} \odot \mathbf{A})^T$, of dimensions $J_{n_B} \times (IK_{n_C})$, for $n_B = 1, \dots, N_B$ and $n_C = 1, \dots, N_C$, and $\mathbf{W}^{(3)}$ can be partitioned into blocks $\mathbf{W}_{n_C,n_B}^{(3)} = \mathbf{C}_{n_C}(\mathbf{B}_{n_B} \odot \mathbf{A})^T$,

of dimensions $K_{n_C} \times (IJ_{n_B})$, for $n_C = 1, \dots, N_C$ and $n_B = 1, \dots, N_B$.⁷

If we partition $\mathbf{X}^{(1)}$, $\mathbf{X}^{(2)}$, and $\mathbf{X}^{(3)}$ accordingly, then we can write

$$\begin{aligned} f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) &= \sum_{n_A=1}^{N_A} \sum_{n_C=1}^{N_C} \frac{1}{2} \|\mathbf{X}_{n_A,n_C}^{(1)} - \mathbf{A}_{n_A}(\mathbf{C}_{n_C} \odot \mathbf{B})^T\|_F^2 \\ &= \sum_{n_B=1}^{N_B} \sum_{n_C=1}^{N_C} \frac{1}{2} \|\mathbf{X}_{n_B,n_C}^{(2)} - \mathbf{B}_{n_B}(\mathbf{C}_{n_C} \odot \mathbf{A})^T\|_F^2 \\ &= \sum_{n_C=1}^{N_C} \sum_{n_B=1}^{N_B} \frac{1}{2} \|\mathbf{X}_{n_C,n_B}^{(3)} - \mathbf{C}_{n_C}(\mathbf{B}_{n_B} \odot \mathbf{A})^T\|_F^2. \end{aligned} \quad (22)$$

These expressions will be fundamental for the development of the distributed ADMoM for large NTF.

B. Distributed ADMoM for Large NTF

In order to put the large NTF problem into ADMoM form, we introduce auxiliary variables $\tilde{\mathbf{A}} = [\tilde{\mathbf{A}}_1^T \ \cdots \ \tilde{\mathbf{A}}_{N_A}^T]^T$, with $\tilde{\mathbf{A}}_{n_A} \in \mathbb{R}^{I_{n_A} \times F}$, for $n_A = 1, \dots, N_A$, $\tilde{\mathbf{B}} = [\tilde{\mathbf{B}}_1^T \ \cdots \ \tilde{\mathbf{B}}_{N_B}^T]^T$, with $\tilde{\mathbf{B}}_{n_B} \in \mathbb{R}^{J_{n_B} \times F}$, for $n_B = 1, \dots, N_B$, and $\tilde{\mathbf{C}} = [\tilde{\mathbf{C}}_1^T \ \cdots \ \tilde{\mathbf{C}}_{N_C}^T]^T$, with $\tilde{\mathbf{C}}_{n_C} \in \mathbb{R}^{K_{n_C} \times F}$, for $n_C = 1, \dots, N_C$, and consider the equivalent problem

$$\begin{aligned} \min_{\mathbf{A}, \tilde{\mathbf{A}}, \mathbf{B}, \tilde{\mathbf{B}}, \mathbf{C}, \tilde{\mathbf{C}}} \quad & f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) + \sum_{n_A=1}^{N_A} g(\tilde{\mathbf{A}}_{n_A}) \\ & + \sum_{n_B=1}^{N_B} g(\tilde{\mathbf{B}}_{n_B}) + \sum_{n_C=1}^{N_C} g(\tilde{\mathbf{C}}_{n_C}) \\ \text{subject to} \quad & \mathbf{A}_{n_A} - \tilde{\mathbf{A}}_{n_A} = \mathbf{0}, \quad n_A = 1, \dots, N_A, \\ & \mathbf{B}_{n_B} - \tilde{\mathbf{B}}_{n_B} = \mathbf{0}, \quad n_B = 1, \dots, N_B, \\ & \mathbf{C}_{n_C} - \tilde{\mathbf{C}}_{n_C} = \mathbf{0}, \quad n_C = 1, \dots, N_C. \end{aligned} \quad (24)$$

If we introduce dual variables $\mathbf{Y}_{\mathbf{A}} = [\mathbf{Y}_{\mathbf{A}_1}^T \ \cdots \ \mathbf{Y}_{\mathbf{A}_{N_A}}^T]^T$, with $\mathbf{Y}_{\mathbf{A}_{n_A}} \in \mathbb{R}^{I_{n_A} \times F}$, for $n_A = 1, \dots, N_A$, $\mathbf{Y}_{\mathbf{B}} = [\mathbf{Y}_{\mathbf{B}_1}^T \ \cdots \ \mathbf{Y}_{\mathbf{B}_{N_B}}^T]^T$, with $\mathbf{Y}_{\mathbf{B}_{n_B}} \in \mathbb{R}^{J_{n_B} \times F}$, for $n_B = 1, \dots, N_B$, and $\mathbf{Y}_{\mathbf{C}} = [\mathbf{Y}_{\mathbf{C}_1}^T \ \cdots \ \mathbf{Y}_{\mathbf{C}_{N_C}}^T]^T$, with $\mathbf{Y}_{\mathbf{C}_{n_C}} \in \mathbb{R}^{K_{n_C} \times F}$, for $n_C = 1, \dots, N_C$, the augmented Lagrangian is written as in (23) at the bottom of the page.

⁷An extension of the above partitioning scheme to higher order tensors appears in Appendix A.

$$\begin{aligned} L_{\rho}(\mathbf{A}, \mathbf{B}, \mathbf{C}, \tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{C}}, \mathbf{Y}_{\mathbf{A}}, \mathbf{Y}_{\mathbf{B}}, \mathbf{Y}_{\mathbf{C}}) &= f_{\underline{\mathbf{X}}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) + \sum_{n_A=1}^{N_A} g(\tilde{\mathbf{A}}_{n_A}) + \sum_{n_B=1}^{N_B} g(\tilde{\mathbf{B}}_{n_B}) + \sum_{n_C=1}^{N_C} g(\tilde{\mathbf{C}}_{n_C}) \\ &+ \sum_{n_A=1}^{N_A} \left(\mathbf{Y}_{\mathbf{A}_{n_A}} * (\mathbf{A}_{n_A} - \tilde{\mathbf{A}}_{n_A}) + \frac{\rho_{\mathbf{A}}}{2} \|\mathbf{A}_{n_A} - \tilde{\mathbf{A}}_{n_A}\|_F^2 \right) \\ &+ \sum_{n_B=1}^{N_B} \left(\mathbf{Y}_{\mathbf{B}_{n_B}} * (\mathbf{B}_{n_B} - \tilde{\mathbf{B}}_{n_B}) + \frac{\rho_{\mathbf{B}}}{2} \|\mathbf{B}_{n_B} - \tilde{\mathbf{B}}_{n_B}\|_F^2 \right) \\ &+ \sum_{n_C=1}^{N_C} \left(\mathbf{Y}_{\mathbf{C}_{n_C}} * (\mathbf{C}_{n_C} - \tilde{\mathbf{C}}_{n_C}) + \frac{\rho_{\mathbf{C}}}{2} \|\mathbf{C}_{n_C} - \tilde{\mathbf{C}}_{n_C}\|_F^2 \right). \end{aligned} \quad (23)$$

The ADMoM for this problem is as follows:

$$\begin{aligned}
(\mathbf{A}^{k+1}, \mathbf{B}^{k+1}, \mathbf{C}^{k+1}) &= \arg \min_{\mathbf{A}, \mathbf{B}, \mathbf{C}} \left(f_{\mathbf{X}}(\mathbf{A}, \mathbf{B}, \mathbf{C}) \right. \\
&+ \sum_{n_A=1}^{N_A} \left(\mathbf{Y}_{\mathbf{A}_{n_A}}^k * \mathbf{A}_{n_A} + \frac{\rho_A}{2} \|\mathbf{A}_{n_A} - \tilde{\mathbf{A}}_{n_A}^k\|_F^2 \right) \\
&+ \sum_{n_B=1}^{N_B} \left(\mathbf{Y}_{\mathbf{B}_{n_B}}^k * \mathbf{B}_{n_B} + \frac{\rho_B}{2} \|\mathbf{B}_{n_B} - \tilde{\mathbf{B}}_{n_B}^k\|_F^2 \right) \\
&+ \left. \sum_{n_C=1}^{N_C} \left(\mathbf{Y}_{\mathbf{C}_{n_C}}^k * \mathbf{C}_{n_C} + \frac{\rho_C}{2} \|\mathbf{C}_{n_C} - \tilde{\mathbf{C}}_{n_C}^k\|_F^2 \right) \right) \\
\tilde{\mathbf{A}}_{n_A}^{k+1} &= \left(\mathbf{A}_{n_A}^{k+1} + \frac{1}{\rho_A} \mathbf{Y}_{\mathbf{A}_{n_A}}^k \right)_+, \quad n_A = 1, \dots, N_A, \\
\tilde{\mathbf{B}}_{n_B}^{k+1} &= \left(\mathbf{B}_{n_B}^{k+1} + \frac{1}{\rho_B} \mathbf{Y}_{\mathbf{B}_{n_B}}^k \right)_+, \quad n_B = 1, \dots, N_B, \\
\tilde{\mathbf{C}}_{n_C}^{k+1} &= \left(\mathbf{C}_{n_C}^{k+1} + \frac{1}{\rho_C} \mathbf{Y}_{\mathbf{C}_{n_C}}^k \right)_+, \quad n_C = 1, \dots, N_C, \\
\mathbf{Y}_{\mathbf{A}_{n_A}}^{k+1} &= \mathbf{Y}_{\mathbf{A}_{n_A}}^k + \rho_A \left(\mathbf{A}_{n_A}^{k+1} - \tilde{\mathbf{A}}_{n_A}^{k+1} \right), \quad n_A = 1, \dots, N_A, \\
\mathbf{Y}_{\mathbf{B}_{n_B}}^{k+1} &= \mathbf{Y}_{\mathbf{B}_{n_B}}^k + \rho_B \left(\mathbf{B}_{n_B}^{k+1} - \tilde{\mathbf{B}}_{n_B}^{k+1} \right), \quad n_B = 1, \dots, N_B, \\
\mathbf{Y}_{\mathbf{C}_{n_C}}^{k+1} &= \mathbf{Y}_{\mathbf{C}_{n_C}}^k + \rho_C \left(\mathbf{C}_{n_C}^{k+1} - \tilde{\mathbf{C}}_{n_C}^{k+1} \right), \quad n_C = 1, \dots, N_C. \quad (25)
\end{aligned}$$

The minimization problem in the first line of (25) is non-convex. Based on (22), we propose the alternating optimization scheme given in (26) at the bottom of the page.

Again, during each ADMoM iteration, we avoid the solution of constrained optimization problems. Furthermore, and more importantly, having computed all algorithm quantities at iteration k , the updates of $\mathbf{A}_{n_A}^k$, for $n_A = 1, \dots, N_A$, are *independent* and can be computed in *parallel*. Then, we can compute in parallel the updates of $\mathbf{B}_{n_B}^k$, for $n_B = 1, \dots, N_B$, and, finally, the updates of $\mathbf{C}_{n_C}^k$, for $n_C = 1, \dots, N_C$.


We note that we can solve problem (24) using the centralized ADMoM of Section IV. In fact, if we initialize the corresponding quantities of the two algorithms with the same values, then the two algorithms evolve in *exactly* the same way. As a result, the study (for example, convergence analysis and/or numerical behavior) of one of them is sufficient for the characterization of both.

Thus, via the distributed ADMoM, we simply uncover the *inherent parallelism* in the updates of the blocks of \mathbf{A}^k , \mathbf{B}^k , and \mathbf{C}^k . In Appendix B, we present a detailed proof of the equivalence of these two forms of ADMoM NTF.

C. A Parallel Implementation of ADMoM for Large NTF

In the sequel, we briefly describe a simple implementation of ADMoM for large NTF on a mesh-type architecture. In order to keep the presentation simple, we assume that (1) $N_A = N_B = N_C = N$ and (2) each of the matrix unfoldings $\mathbf{X}^{(1)}$, $\mathbf{X}^{(2)}$, and $\mathbf{X}^{(3)}$ has been split into N^2 blocks, with their (i, j) -th blocks stored at the (i, j) -th processing element, for $i, j = 1, \dots, N$ (for related results in the matrix factorization context see [34]).

$$\begin{aligned}
\mathbf{A}_{n_A}^{k+1} &= \arg \min_{\mathbf{A}_{n_A}} \left(\left(\sum_{n_C=1}^{N_C} \frac{1}{2} \|\mathbf{X}_{n_A, n_C}^{(1)} - \mathbf{A}_{n_A} (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k)^T\|_F^2 \right) + \mathbf{Y}_{\mathbf{A}_{n_A}}^k * \mathbf{A}_{n_A} + \frac{\rho_A}{2} \|\mathbf{A}_{n_A} - \tilde{\mathbf{A}}_{n_A}^k\|_F^2 \right) \\
&= \left(\left(\sum_{n_C=1}^{N_C} \mathbf{X}_{n_A, n_C}^{(1)} (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k) \right) + \rho_A \tilde{\mathbf{A}}_{n_A}^k - \mathbf{Y}_{\mathbf{A}_{n_A}}^k \right) \\
&\quad \left(\left(\sum_{n_C=1}^{N_C} (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k)^T (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k) \right) + \rho_A \mathbf{I}_F \right)^{-1}, \quad \text{for } n_A = 1, \dots, N_A, \\
\mathbf{B}_{n_B}^{k+1} &= \arg \min_{\mathbf{B}_{n_B}} \left(\left(\sum_{n_C=1}^{N_C} \frac{1}{2} \|\mathbf{X}_{n_B, n_C}^{(2)} - \mathbf{B}_{n_B} (\mathbf{C}_{n_C}^k \odot \mathbf{A}^{k+1})^T\|_F^2 \right) + \mathbf{Y}_{\mathbf{B}_{n_B}}^k * \mathbf{B}_{n_B} + \frac{\rho_B}{2} \|\mathbf{B}_{n_B} - \tilde{\mathbf{B}}_{n_B}^k\|_F^2 \right) \\
&= \left(\left(\sum_{n_C=1}^{N_C} \mathbf{X}_{n_B, n_C}^{(2)} (\mathbf{C}_{n_C}^k \odot \mathbf{A}^{k+1}) \right) + \rho_B \tilde{\mathbf{B}}_{n_B}^k - \mathbf{Y}_{\mathbf{B}_{n_B}}^k \right) \\
&\quad \left(\left(\sum_{n_C=1}^{N_C} (\mathbf{C}_{n_C}^k \odot \mathbf{A}^{k+1})^T (\mathbf{C}_{n_C}^k \odot \mathbf{A}^{k+1}) \right) + \rho_B \mathbf{I}_F \right)^{-1}, \quad \text{for } n_B = 1, \dots, N_B, \\
\mathbf{C}_{n_C}^{k+1} &= \arg \min_{\mathbf{C}_{n_C}} \left(\left(\sum_{n_B=1}^{N_B} \frac{1}{2} \|\mathbf{X}_{n_C, n_B}^{(3)} - \mathbf{C}_{n_C} (\mathbf{B}_{n_B}^{k+1} \odot \mathbf{A}^{k+1})^T\|_F^2 \right) + \mathbf{Y}_{\mathbf{C}_{n_C}}^k * \mathbf{C}_{n_C} + \frac{\rho_C}{2} \|\mathbf{C}_{n_C} - \tilde{\mathbf{C}}_{n_C}^k\|_F^2 \right) \\
&= \left(\left(\sum_{n_B=1}^{N_B} \mathbf{X}_{n_C, n_B}^{(3)} (\mathbf{B}_{n_B}^{k+1} \odot \mathbf{A}^{k+1}) \right) + \rho_C \tilde{\mathbf{C}}_{n_C}^k - \mathbf{Y}_{\mathbf{C}_{n_C}}^k \right) \\
&\quad \left(\left(\sum_{n_B=1}^{N_B} (\mathbf{B}_{n_B}^{k+1} \odot \mathbf{A}^{k+1})^T (\mathbf{B}_{n_B}^{k+1} \odot \mathbf{A}^{k+1}) \right) + \rho_C \mathbf{I}_F \right)^{-1}, \quad \text{for } n_C = 1, \dots, N_C. \quad (26)
\end{aligned}$$

Fig. 1. Distributed computation of \mathbf{A}_n^{k+1} , for $n = 1, \dots, N$.

In Fig. 1, we depict the data flow for the computation of the blocks of \mathbf{A}^{k+1} . The inputs to the N top processing elements are \mathbf{C}_n^k , for $n = 1, \dots, N$, as well as \mathbf{B}^k , which is common input to all top processing elements. Each processing element uses its inputs and memory contents and computes certain partial matrix sums. The communications between the processing elements are local and involve either the forwarding of the terms \mathbf{C}_n^k , for $n = 1, \dots, N$, and \mathbf{B}^k (top-down communication), or the forwarding of the partial sums $\sum_{l=1}^j \mathbf{X}_{n,l}^{(1)}(\mathbf{C}_l^k \odot \mathbf{B}^k)$ and $\sum_{l=1}^j (\mathbf{C}_l^k \odot \mathbf{B}^k)^T (\mathbf{C}_l^k \odot \mathbf{B}^k)$ (left-right communication), of dimensions $\frac{I}{N} \times F$ and $F \times F$, respectively. The computation of \mathbf{A}_n^{k+1} , for $n = 1, \dots, N$, amounts to solution of ρ systems of linear equations with common coefficient matrix and takes place at the rightmost computing elements.

Then, using a similar strategy, we can compute the blocks of \mathbf{B}^{k+1} and, finally, the blocks of \mathbf{C}^{k+1} . The updates of the auxiliary and dual variables are very simple and can be performed locally (see at the rightmost computing elements of Fig. 1).

As we see in Fig. 1, in order to compute the blocks of the \mathbf{A}^{k+1} , we use the appropriate blocks of \mathbf{C}_n^k , for $n = 1, \dots, N$, as well as the *whole* matrix \mathbf{B}^k . When the size of \mathbf{B}^k is not very large, the communication cost is *not* prohibitive (analogous arguments holds for the computation of the blocks of \mathbf{B}^{k+1} and \mathbf{C}^{k+1}). Of course, if one or more latent factors are very large, the communication cost significantly increases.

Concerning the distributed implementation of ADMoM for large tensor factorization with structural constraints other than the non-negativity of the latent factors, we note that, if the size of the latent factors is not very large, the communication cost of gathering together the blocks of the auxiliary variables $\tilde{\mathbf{A}}^k$, $\tilde{\mathbf{B}}^k$, and $\tilde{\mathbf{C}}^k$, is not prohibitive, enabling the computation of more complicated non-separable projections, like, for example, projection onto the set of non-negative matrices with a certain maximum number of non-negative elements.

Actual implementation of the distributed ADMoM for large NTF will depend on the specific parallel architecture and programming environment used. Since our aim in this paper is to introduce the basic methodology and computational framework, we leave those customizations and performance tune-ups, which are further away from the signal processing core, for follow-up work to be reported in the high-performance computing literature.

VI. NUMERICAL EXPERIMENTS

A. Comparison of ADMoM With NALS and NLS

In our numerical experiments, we compare ADMoM NTF with (1) NALS NTF, as implemented in the `parafac` routine of the *N-way toolbox* for Matlab [35] and (2) NTF using the nonlinear least-squares solvers (NLS), as implemented in the `sdf_nls` routine of `tensorlab` [36] (with the non-negativity option turned on in both cases). In all cases, we use random initialization. More specifically, the initialization of the ADMoM NTF is as follows. We give non-negative *random* values to \mathbf{B}^0 and \mathbf{C}^0 and zero values to the other state variables of the algorithm, namely, $\tilde{\mathbf{A}}^0$, $\tilde{\mathbf{B}}^0$, $\tilde{\mathbf{C}}^0$, $\mathbf{Y}_{\mathbf{A}}^0$, $\mathbf{Y}_{\mathbf{B}}^0$, and $\mathbf{Y}_{\mathbf{C}}^0$.⁸

In extensive numerical experiments, we have observed that the relative performance of the algorithms depends on the size and rank of the tensor as well as the additive noise power. Thus, we consider 12 different scenarios, corresponding to the combinations of the following cases:

- 1) one, two, or three tensor dimensions are large;
- 2) rank F is small or large;
- 3) additive noise is weak or strong.

⁸In certain cases, it may be possible to employ algebraic initialization schemes (e.g., see [16] and references therein), but for NTF we observed that they perform (very) well only in (very) high SNR cases. In the same vein, projecting an unconstrained decomposition to force non-negativity *a posteriori*, only seems to work in high SNR cases.

For each scenario, we generate $R = 50$ realizations of tensor $\underline{\mathbf{X}}$ as follows. We generate random matrices \mathbf{A}^o , \mathbf{B}^o , and \mathbf{C}^o with i.i.d. $\mathcal{U}[0, 1]$ elements (using the `rand` command of Matlab) and construct $\underline{\mathbf{X}} = [\mathbf{A}^o, \mathbf{B}^o, \mathbf{C}^o] + \underline{\mathbf{N}}$, where $\underline{\mathbf{N}}$ consists of i.i.d. $\mathcal{N}(0, \sigma_N^2)$ elements. For each realization, we solve the NTF problem with (1) NALS (`parafac`), (2) NLS (`sdf_nls`), and (3) ADMoM.

We designed our experiments so that, upon convergence, all algorithms achieve practically the same relative factorization error. Towards this end, we set the values of the stopping parameters as follows: the parameter `Options(1)` of `parafac` is set to `Options(1) = 10-5`, the parameter `TolFun` of `sdf_nls` is set to `TolFun = 10-8`, and the ADMoM stopping parameters are set to $\epsilon^{\text{abs}} = 10^{-4}$ and $\epsilon^{\text{rel}} = 10^{-4}$.

In all cases, the initial values of the ADMoM penalty terms are $\rho_{\mathbf{M}} = 1$, for $\mathbf{M} = \mathbf{A}, \mathbf{B}, \mathbf{C}$, while the ADMoM penalty term adaptation parameters are $\mu = 8$, $\tau^{\text{incr}} = 4$, $\tau^{\text{decr}} = 2$.

In practice, convergence properties of ADMoM NTF depend on the (random) initialization point. In some cases, convergence may be quite fast while, in others, it may be quite slow. As we shall see in the sequel, this phenomenon seems more prominent in the cases where rank F is large. In order to overcome the slow convergence properties associated with bad initial points, we adopted the following strategy. We execute ADMoM NTF for up to $n_{\text{max}} = 400$ iterations (we have observed that, in the great majority of the cases in the scenarios we examined, this number of iterations is sufficient for convergence when we start from a good initial point). If ADMoM does not converge within n_{max} iterations, then we restart it from another random initial point; we repeat this procedure until ADMoM converges.⁹

Before proceeding, we mention that *all* the algorithms converged in *all* the realizations we run.

Since an accurate statement about the computational complexity per iteration of `parafac` is not easy, the metric we used for comparison of the algorithms is the `cputime` of Matlab. Despite the fact that `cputime` is strongly dependent on the computer hardware and the actual algorithm implementation, we feel that it is a useful metric for the assessment of the relative efficiency of the algorithms.¹⁰ The reason is that we used carefully developed, publicly available Matlab toolbox implementations of the baseline algorithms, and we carefully coded our ADMoM NTF implementation.

In Table I, we present the mean and standard deviation of `cputime`, in seconds, denoted as `mean(t)` and `std(t)`, respectively, for NALS, NLS, and ADMoM. We also present the mean relative factorization error (which is common to all algorithms up to four decimal digits), defined as

$$\text{mean(RFE)} := \frac{1}{n} \sum_{k=1}^n \frac{\|\underline{\mathbf{X}}_k - [\mathbf{A}_k, \mathbf{B}_k, \mathbf{C}_k]\|_F}{\|\underline{\mathbf{X}}_k\|_F},$$

where $\underline{\mathbf{X}}_k$ is the k -th noisy tensor realization and \mathbf{A}_k , \mathbf{B}_k , and \mathbf{C}_k are the factors returned by a factorization algorithm. Our observations are as follows:

- 1) There is no clear winner. Certainly, for high ranks, NLS has very good behavior.
- 2) In general, both NALS and NLS have more predictable behavior than ADMoM. Especially for high ranks, the




Fig. 2. `cputime` for $I = 3000$, $J = K = 50$, $F = 3$, and $\sigma_N^2 = 10^{-2}$. NALS (blue solid line), NLS (green dashed line), ADMoM (red dotted-dashed line).




Fig. 3. `cputime` for $I = 3000$, $J = K = 50$, $F = 30$ and $\sigma_N^2 = 10^{-2}$. NALS (blue solid line), NLS (green dashed line), ADMoM (red dotted-dashed line).

`cputime` of our implementation of ADMoM has large variance.

- 3) For small ranks, ADMoM looks more competitive and, in the cases where one dimension is much larger than the other two, it behaves very well (we shall say more on this later).

In order to get a better feeling of the behavior of the three algorithms, we plot their `cputime`, along the 50 realizations we used to obtain the averages of Table I, for two different scenarios. In Fig. 2, we consider the case for $I = 3000$, $J = K = 50$, $F = 3$ and $\sigma_N^2 = 10^{-2}$. We observe that the behavior of the algorithms is stable, in the sense that there is a clear ordering among the three algorithms, with no large variations. In Fig. 3, we keep the dimensions and the noise power the same as before

⁹Of course, one may think of more elaborate strategies such as, for example, running in parallel more than one versions of the algorithm, with different initializations.

¹⁰For our experiments, we run Matlab 2014a on a MacBook Pro with a 2.5 GHz Intel Core i7 Intel processor and 16 GB RAM.

TABLE I
MEAN RELATIVE FACTORIZATION ERROR AND MEAN AND STANDARD DEVIATION OF **cputime**, IN SEC, FOR NALS, NLS AND ADMoM NTF

Size	F	σ_N^2	mean(RFE)	NALS mean(t)	NLS mean(t)	ADMoM mean(t)	NALS std(t)	NLS std(t)	ADMoM std(t)
$3000 \times 50 \times 50$	3	10^{-2}	0.2156	12.4010	17.2584	7.1806	1.4770	6.9348	3.8647
		10^{-4}	0.0221	16.6500	16.5962	7.5098	1.9625	3.3589	4.2499
	30	10^{-2}	0.0260	212.0598	115.7030	110.4128	11.4579	8.0409	91.4882
		10^{-4}	0.0026	270.6674	117.2152	148.5052	11.5324	11.3154	149.0542
$400 \times 400 \times 50$	3	10^{-2}	0.2175	8.6174	4.7124	8.0710	0.9264	1.4596	4.4952
		10^{-4}	0.0222	11.0670	4.8916	8.3614	1.2075	1.6443	3.1833
	30	10^{-2}	0.0260	71.1950	31.7546	106.1362	8.2119	3.2162	76.5743
		10^{-4}	0.0026	92.4838	31.1690	94.5734	7.9280	3.7062	88.1671
$200 \times 200 \times 200$	5	10^{-2}	0.1400	10.9142	4.1756	12.5190	1.2783	0.7915	2.1334
		10^{-4}	0.0143	14.2882	4.1070	12.6184	2.5817	0.9059	2.1616
	30	10^{-2}	0.0260	55.0806	16.1838	33.9268	4.4886	1.3649	12.2687
		10^{-4}	0.0026	70.0238	16.7624	32.1670	6.5737	1.4152	10.0182

TABLE II
MEAN RELATIVE FACTORIZATION ERROR AND MEAN AND STANDARD DEVIATION OF **cputime**, IN SEC, FOR NALS, NLS AND ADMoM NTF FOR $I = 10^4$, $J = K = 50$, $F = 10$, AND $\sigma_N^2 = 10^{-2}$.

NALS mean(RFE)	NLS mean(RFE)	ADMoM mean(RFE)	NALS mean(t)	NLS mean(t)	ADMoM mean(t)	NALS std(t)	NLS std(t)	ADMoM std(t)
0.0752	0.0758	0.0751	116.0722	201.6556	30.6460	19.8483	38.5401	16.5823

and increase the rank to $F = 30$. We observe that the variance of ADMoM **cputime** has significantly increased, while both NALS and NLS show stable behavior. When ADMoM starts from a good initial point, it converges faster than NALS and NLS while, when it starts from bad initial points, it needs one or more restarts.

In order to check if ADMoM maintains its advantage over NALS and NLS in the cases where one dimension is very large, compared with the other two, and the rank is relatively small, we performed an experiment with $I = 10^4$, $J = K = 50$, $F = 10$, and $\sigma_N^2 = 10^{-2}$. However, in this case, we used somewhat relaxed stopping conditions for all algorithms; more specifically, we used $\text{Options}(1) = 10^{-3}$, $\text{TolFun} = 10^{-6}$, $\epsilon^{\text{abs}} = 10^{-3}$, and $\epsilon^{\text{rel}} = 10^{-3}$. In Table II, we present the mean relative factorization errors and the mean and standard deviation of **cputime**. As we can see, both NALS and NLS are slightly less accurate than ADMoM, in terms of relative factorization error, which means that their stopping criteria are more relaxed. In terms of **cputime**, we see that ADMoM is much faster than both NALS and NLS. In Fig. 4, we plot the **cputime** of the three algorithms for the 50 realizations of the experiment. Again, we see the significant difference between ADMoM and both NALS and NLS. We note that if we had used as values of the stopping parameters those of our initial experiments, then the gain of ADMoM, compared with NALS and NLS, would have been much greater. However, we believe that we have made clear that, in this case, ADMoM has a clear advantage. We have made analogous observations for larger I .

B. A Closer Look at ADMoM

In order to get a more detailed view of the convergence properties of ADMoM, we return to the scenario with $I = 3000$, $J = K = 50$, $F = 30$, and $\sigma_N^2 = 10^{-2}$, whose **cputime** we plot in Fig. 3. We recall that, in order to converge in this case, ADMoM needed often restarts. In Fig. 5, we plot the total number of ADMoM iterations, denoted as **iters**, and the number of ADMoM iterations during its final way to




Fig. 4. **cputime** for $I = 10^4$, $J = K = 50$, $F = 10$, and $\sigma_N^2 = 10^{-2}$. NALS (blue solid line), NLS (green dashed line), ADMoM (red dotted-dashed line).




Fig. 5. Number of ADMoM iterations for $I = 3000$, $J = 50 = K = 50$, $F = 30$, and $\sigma_N^2 = 10^{-2}$; **iters** (blue line), and $\text{mod}(\text{iters}, n_{\text{max}})$ (green line).

convergence, which is equal to $\text{mod}(\text{iters}, n_{\max})$. As expected, iters is compatible with the corresponding cputime (see the red line in Fig. 3). Quantity $\text{mod}(\text{iters}, n_{\max})$ shows how many iterations are required for convergence if ADMoM always starts from good initial points. We observe that $\text{mod}(\text{iters}, n_{\max})$ is quite stable around its mean, which is approximately equal to 320. This gives an estimate of the fastest possible ADMoM convergence in this case.

C. ADMoM NTF With Under- and Over-Estimated Rank

In the sequel, we consider ADMoM behavior in the cases where we under- or over-estimate the true rank, in both noisy and noiseless cases. Towards this end, we fix $I = J = K = 100$ and $F = 30$, and investigate ADMoM with exact rank as well as with rank under- and over-estimated by 1. We expect that, in this case, all versions of ADMoM may need restarts. In the sequel, we examine the influence of under- and over-estimating the rank on (1) factorization accuracy and (2) number of restarts. Of course, in under-modeled cases, we expect that the relative factorization error will be higher than that of the true rank case. However, we know nothing in advance about ADMoM behavior in over-modeled cases. In order to get insight into these issues, we perform the following experiment. We set stopping parameters $\epsilon^{\text{abs}} = \epsilon^{\text{rel}} = 10^{-4}$ and run each of the three versions of ADMoM for $n_{\max} = 500$ iterations. For each version, we proceed as follows: if it converges within n_{\max} iterations, we stop; otherwise, we restart, and repeat until convergence. Thus, finally, the number of iterations for each ADMoM version will be a multiple of n_{\max} . For the computation of the trajectory of the mean relative factorization error we use only the last n_{\max} values; in this way, we avoid the influence of bad initial points. However, we keep count of the restarts of each version and, thus, can assess the time it needs to achieve convergence.

In Fig. 6, we plot the average relative factorization errors (computed over 50 realizations in the way we mentioned before), versus the iteration number, for $\sigma_N^2 = 10^{-2}$. As was expected, the ADMoM version with under-estimated rank converges to a higher relative factorization error. We observe that the average relative factorization errors for ADMoM with exact rank and rank over-estimated by 1 follow almost the same trajectory. The average numbers of restarts for the three ADMoM versions are $\text{mean}(\text{restarts}_{\text{exact}}) = 1.08$, $\text{mean}(\text{restarts}_{\text{under}}) = 1.22$, and $\text{mean}(\text{restarts}_{\text{over}}) = 1.92$. Thus, in the cases of over-estimated rank, we finally achieve a relative factorization error trajectory as good as in the exact rank case, but we may need more restarts and, thus, more time. This implies that the probability of bad initial points may increase.

In Fig. 7, we plot the same quantities for noiseless data. Again, the ADMoM behavior in the under-modeled case is as expected. Interestingly, we observe that there is no relative factorization error floor neither for the exact rank nor for the over-estimated by 1 rank case. Reasonably, after a certain precision level, the over-modeled case converges slower. The average numbers of restarts for the three ADMoM versions are $\text{mean}(\text{restarts}_{\text{exact}}) = 1.14$, $\text{mean}(\text{restarts}_{\text{under}}) = 1.54$, and $\text{mean}(\text{restarts}_{\text{over}}) = 1.06$.




Fig. 6. Average relative factorization errors for $I = J = K = 100$, $F = 30$, and $\sigma_N^2 = 0.1$. Exact rank case (solid blue line), over-estimated rank by 1 (dotted-dashed red line), underestimated rank by 1 (green dashed line).




Fig. 7. Average relative factorization errors for noiseless case with $I = J = K = 100$ and $F = 30$. Exact rank case (solid blue line), over-estimated rank by 1 (dotted-dashed red line), underestimated rank by 1 (green dashed line).

We have observed similar behavior for more drastic rank under- and over-estimation.

D. ADMoM for NTF With Box-Linear Constraints

In our final experiment, we briefly consider NTF for the case where two of the latent factors, say \mathbf{A} and \mathbf{B} , are non-negative while \mathbf{C} is subject to box-linear constraints in the sense that each row of \mathbf{C} is a probability mass function, that is, has non-negative elements with sum equal to 1.

The only difference between the ADMoM for this case and the ADMoM for NTF is that, instead of computing \mathbf{C}^{k+1} as the solution of an unconstrained least-squares problem, we compute it as the solution of linearly constrained least-squares; note that both cases exhibit closed-form solutions.

In Fig. 8, we illustrate the behavior of ADMoM in this case by plotting the trajectories of the norms of the average (over 50 realizations) relative estimation errors of the latent factors, as computed by function `cpderr` of `tensorlab`, versus the iteration number, for a noiseless case with $I = J = 100$, $K = 50$ and $F = 5$. We observe that ADMoM works to very high precision.




Fig. 8. Average relative estimation errors for noiseless case with $I = J = 100$, $K = 50$ and $F = 20$. Factors **A** (solid blue line), **B** (dashed red line), and **C** (green dotted-dashed line).

E. Discussion

Our numerical results are encouraging and suggest that, in many cases, ADMoM NTF can efficiently achieve close to state-of-the-art factorization accuracy. The fact that ADMoM is suitable for high-performance parallel implementation (the first NTF algorithm with this property, as far as we know) can only increase its potential. Thus, we believe that it will be a valuable tool in the NTF toolbox.

Obviously, in order to fully uncover the pros and cons of ADMoM NTF, more extensive experimentation is required. But our intention in this paper is to give the fundamental ideas and some basic performance metrics. Experiments with real-world data (using ADMoM for tensor completion and factorization) as well as constraints well beyond non-negativity are ongoing work.

A weak point of the version of ADMoM we developed in this manuscript is the high `cputime` variance in cases of high rank, where ADMoM seems more sensitive than other algorithms with respect to initialization. The improvement of the ADMoM behavior in these cases remains a very interesting problem. To achieve this goal, it might be possible to combine elements of NLS and ADMoM and derive a more efficient algorithm. However, more research efforts are needed in this direction.

As we mentioned, if the centralized and the distributed algorithms start from the same initial point, they evolve in exactly the same way. Thus, distributed ADMoM inherits the convergence properties of centralized ADMoM.

VII. CONCLUSION

Motivated by emerging big data applications, involving multi-way tensor data, and the ensuing need for scalable tensor factorization tools, we developed a new constrained tensor factorization framework based on the ADMoM. We used non-negative factorization of third order tensors as an example to work out the main ideas, but our approach can be generalized to higher order tensors, many other types of constraints on the latent factors, as well as other tensor factorizations and tensor

completion. Our numerical experiments were encouraging, indicating that, in many cases, the ADMoM-based NTF has high potential as an alternative to the state-of-the-art and, in some cases, it may become state-of-the-art. The fact that it is naturally amenable to parallel implementation can only increase its potential. The improvement of its behavior in the high rank cases remains a very interesting problem.

APPENDIX A

EXTENSION TO HIGHER ORDER TENSORS

In this appendix, we highlight how our approach can be extended to higher order tensors. We focus on fourth-order tensors, with the general case being obvious. If $\underline{\mathbf{W}} = [\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D}]$, then its matrix unfoldings satisfy relations

$$\begin{aligned}\mathbf{W}^{(1)} &= \mathbf{A}(\mathbf{D} \odot (\mathbf{C} \odot \mathbf{B}))^T, \\ \mathbf{W}^{(2)} &= \mathbf{B}(\mathbf{D} \odot (\mathbf{C} \odot \mathbf{A}))^T, \\ \mathbf{W}^{(3)} &= \mathbf{C}(\mathbf{D} \odot (\mathbf{B} \odot \mathbf{A}))^T, \\ \mathbf{W}^{(4)} &= \mathbf{D}(\mathbf{C} \odot (\mathbf{B} \odot \mathbf{A}))^T.\end{aligned}$$

Partitioning matrices **A**, **B**, **C**, and **D** as in Section V-A, we obtain that matrix $\mathbf{W}^{(1)}$ can be partitioned into $N_A \times N_D$ blocks, with the (i, j) -th block being equal to

$$\mathbf{W}_{i,j}^{(1)} = \mathbf{A}_i(\mathbf{D}_j \odot (\mathbf{C} \odot \mathbf{B}))^T.$$

Analogous partitionings apply to the other matrix unfoldings. Then, development of ADMoM NTF (centralized and distributed) is rather easy.

APPENDIX B

ON THE EQUIVALENCE OF THE CENTRALIZED AND THE DISTRIBUTED ADMoM NTF

A simple proof of the equivalence of the centralized and the distributed ADMoM NTF is as follows. We focus on the update of \mathbf{A}^k of the centralized algorithm and the updates of its blocks, $\mathbf{A}_{n_A}^k$, for $n_A = 1, \dots, N_A$, of the distributed algorithm, and prove that they are equivalent. We remind that

$$\begin{aligned}\mathbf{A}^{k+1} &= \left(\mathbf{X}^{(1)}(\mathbf{C}^k \odot \mathbf{B}^k) + \rho_{\mathbf{A}} \tilde{\mathbf{A}}^k - \mathbf{Y}_{\mathbf{A}}^k \right) \\ &\quad \left((\mathbf{C}^k \odot \mathbf{B}^k)^T (\mathbf{C}^k \odot \mathbf{B}^k) + \rho_{\mathbf{A}} \mathbf{I}_F \right)^{-1}.\end{aligned}$$

Using the partitionings of $\mathbf{C}^k \odot \mathbf{B}^k$ (see Section V-A), it can be shown that

$$(\mathbf{C}^k \odot \mathbf{B}^k)^T (\mathbf{C}^k \odot \mathbf{B}^k) = \sum_{n_C=1}^{N_C} (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k)^T (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k).$$

Rewriting the update of \mathbf{A}^k in terms of partitioned matrices, we obtain (27) at the top of the next page. If we focus on a certain block of \mathbf{A}^{k+1} in (27), then we obtain the corresponding update of the distributed algorithm (see (26)). We observe that the matrix inverse in the second line of (27) is *common* to all blocks, and should be computed once.

$$\begin{aligned}
\begin{bmatrix} \mathbf{A}_1^{k+1} \\ \vdots \\ \mathbf{A}_{N_A}^{k+1} \end{bmatrix} &= \left(\begin{bmatrix} \mathbf{X}_{1,1}^{(1)} & \cdots & \mathbf{X}_{1,N_C}^{(1)} \\ \vdots & \ddots & \vdots \\ \mathbf{X}_{N_A,1}^{(1)} & \cdots & \mathbf{X}_{N_A,N_C}^{(1)} \end{bmatrix} \begin{bmatrix} \mathbf{C}_1^k \odot \mathbf{B}^k \\ \vdots \\ \mathbf{C}_{N_C}^k \odot \mathbf{B}^k \end{bmatrix} + \begin{bmatrix} \rho_{\mathbf{A}} \tilde{\mathbf{A}}_1^k - \mathbf{Y}_{\mathbf{A}_1}^k \\ \vdots \\ \rho_{\mathbf{A}} \tilde{\mathbf{A}}_{N_A}^k - \mathbf{Y}_{\mathbf{A}_{N_A}}^k \end{bmatrix} \right) \\
&\quad \times \left(\left(\sum_{n_C=1}^{N_C} (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k)^T (\mathbf{C}_{n_C}^k \odot \mathbf{B}^k) \right) + \rho_{\mathbf{A}} \mathbf{I}_F \right)^{-1}. \quad (27)
\end{aligned}$$

Analogous statements hold for the updates of \mathbf{B}^k and \mathbf{C}^k . The equivalence of the updates of the rest of the variables is trivial.

Thus, in fact, using the partitionings of Section V-A, the distributed ADMoM simply uncovered the inherent parallelism of the centralized ADMoM.

REFERENCES

- [1] A. P. Liavas and N. D. Sidiropoulos, "Parallel algorithms for large-scale constrained tensor decomposition," presented at the IEEE ICASSP, Brisbane, Australia, Apr. 19–24, 2015.
- [2] N. D. Sidiropoulos, R. Bro, and G. Giannakis, "Parallel factor analysis in sensor array processing," *IEEE Trans. Signal Process.*, vol. 48, no. 8, pp. 2377–2388, 2000.
- [3] N. D. Sidiropoulos, G. Giannakis, and R. Bro, "Blind PARAFAC receivers for DS-CDMA systems," *IEEE Trans. Signal Process.*, vol. 48, no. 3, pp. 810–823, 2000.
- [4] D. Nion, K. Mokios, N. D. Sidiropoulos, and A. Potamianos, "Batch and adaptive PARAFAC-based blind separation of convolutive speech mixtures," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 18, no. 6, pp. 1193–1207, 2010.
- [5] C. Faveotte and A. Ozerov, "Notes on nonnegative tensor factorization of the spectrogram for audio source separation: Statistical insights and towards self-clustering of the spatial cues," in *Exploring Music Contents*, ser. Lecture Notes in Computer Science, S. Ystad, M. Aramaki, R. Kronland-Martinet, and K. Jensen, Eds. Berlin, Germany: Springer, 2011, vol. 6684, pp. 102–115.
- [6] E. E. Papalexakis, C. Faloutsos, and N. D. Sidiropoulos, "Parcube: Sparse parallelizable tensor decompositions," in *ECML/PKDD (I)*, ser. Lecture Notes in Computer Science, P. A. Flach, T. D. Bie, and N. Cristianini, Eds. New York, NY, USA: Springer, 2012, vol. 7523, pp. 521–536.
- [7] R. Bro and N. D. Sidiropoulos, "Least squares regression under unimodality and non-negativity constraints," *J. Chemometrics*, vol. 12, pp. 223–247, 1998.
- [8] A. Cichocki, D. Mandic, C. Caiafa, A.-H. Phan, G. Zhou, Q. Zhao, and L. De Lathauwer, "Tensor decompositions for signal processing applications: From two-way to multiway component analysis," *IEEE Signal Process. Mag.*, vol. 32, no. 2, pp. 145–163, Feb. 2014.
- [9] R. Harshman, "Foundations of the PARAFAC procedure: Models and conditions for an "explanatory" multimodal factor analysis," *UCLA Working Papers in Phonetics*, vol. 16, pp. 1–84, 1970.
- [10] R. Harshman, "Determination and proof of minimum uniqueness conditions for PARAFAC-1," *UCLA Working Papers in Phonetics*, vol. 22, pp. 111–117, 1972.
- [11] J. Carroll and 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.
- [12] L. Tucker, "Some mathematical notes on three-mode factor analysis," *Psychometrika*, vol. 31, no. 3, pp. 279–311, Sep. 1966.
- [13] L. De Lathauwer, "Decompositions of a higher-order tensor in block terms—Part II: Definitions and uniqueness," *SIAM J. Matrix Anal. Appl.*, vol. 30, no. 3, pp. 1033–1066, 2008.
- [14] A. Smilde, R. Bro, P. Geladi, and J. Wiley, *Multi-Way Analysis With Applications in the Chemical Sciences*. New York, NY, USA: Wiley, 2004.
- [15] C. Hillar and L.-H. Lim, "Most tensor problems are NP-hard," 2009 [Online]. Available: <http://arxiv.org/abs/0911.1393>
- [16] G. Tomasi and R. Bro, "A comparison of algorithms for fitting the PARAFAC model," *Comput. Statist. Data Anal.*, vol. 50, no. 7, pp. 1700–1734, 2006.
- [17] E. Acar, D. M. Dunlavy, and T. G. Kolda, "A scalable optimization approach for fitting canonical tensor decompositions," *J. Chemometrics* vol. 25, no. 2, pp. 67–86, 2011 [Online]. Available: <http://dx.doi.org/10.1002/cem.1335>
- [18] L. Sorber, M. van Barel, and L. De Lathauwer, "Optimization-based algorithms for tensor decompositions: Canonical polyadic decomposition, decomposition in rank- $(l_r, l_r, 1)$ terms, and a new generalization," *SIAM J. Optim.*, vol. 23, no. 2, pp. 695–720, 2013.
- [19] Hadoop Apache [Online]. Available: <http://hadoop.apache.org/>
- [20] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Commun. ACM*, vol. 51, no. 1, pp. 107–113, 2008.
- [21] U. Kang, E. E. Papalexakis, A. Harpale, and C. Faloutsos, "Gigatensor: scaling tensor analysis up by 100 times—algorithms and discoveries," in *Proc. 18th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min.*, 2012, pp. 316–324.
- [22] N. D. Sidiropoulos, E. E. Papalexakis, and C. Faloutsos, "A parallel algorithm for big tensor decomposition using randomly compressed cubes (PARACOMP)," presented at the IEEE ICASSP, Florence, Italy, May 4–9, 2014.
- [23] N. D. Sidiropoulos, E. E. Papalexakis, and C. Faloutsos, "Parallel randomly compressed cubes: A scalable distributed architecture for big tensor decomposition," *IEEE Signal Process. Mag.*, vol. 31, no. 5, pp. 57–70, Sep. 2014.
- [24] A. de Almeida and A. Kibangou, "Distributed computation of tensor decompositions in collaborative networks," in *Proc. IEEE 5th Int. Workshop Comput. Adv. Multi-Sensor Adapt. Process. (CAMSAP)*, Dec. 2013, pp. 232–235.
- [25] A. de Almeida and A. Kibangou, "Distributed large-scale tensor decomposition," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process. (ICASSP)*, May 2014, pp. 26–30.
- [26] 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 [Online]. Available: <http://dx.doi.org/10.1561/22000000016>
- [27] B. W. Bader and T. G. Kolda, "Efficient MATLAB computations with sparse and factored tensors," *SIAM J. Scientif. Comput.* vol. 30, no. 1, pp. 205–231, 2008 [Online]. Available: <http://dx.doi.org/10.1137/060676489>
- [28] T. G. Kolda and J. Sun, "Scalable tensor decompositions for multi-aspect data mining," in *Proc. IEEE ICDM*, 2008, pp. 363–372.
- [29] B. W. Bader et al., Matlab Tensor Toolbox. ver. 2.5, Jan. 2012 [Online]. Available: <http://www.sandia.gov/tgkolda/TensorToolbox/>
- [30] N. Ravindran, N. D. Sidiropoulos, S. Smith, and G. Karypis, "Memory-efficient parallel computation of tensor and matrix products for big tensor decomposition," presented at the Asilomar Conf. Signals, Syst., Comput., Nov. 3–5, 2014.
- [31] Y. Xu, W. Yin, Z. Wen, and Y. Zhang, "An alternating direction algorithm for matrix completion with nonnegative factors," *Frontiers Math. China*, vol. 51, no. 2, pp. 365–384, 2010.
- [32] A. P. Liavas and N. D. Sidiropoulos, "Parallel algorithms for constrained tensor factorization via the alternating direction method of multipliers," Tech. Univ. of Crete, Greece, Tech. Rep., 2014.
- [33] L. Xu, B. Yu, and Y. Zhang, "An alternating direction and projection algorithm for structure-enforced matrix factorization," 2013 [Online]. Available: <http://www.caam.rice.edu/yzhang/reports/tr1311.pdf>
- [34] R. Gemulla, E. Nijkamp, P. J. Haas, and Y. Sismanis, "Large-scale matrix factorization with distributed stochastic gradient descent," in *Proc. 17th ACM SIGKDD Int. Conf. Knowl. Discov. Data Min. (KDD'11)*, New York, NY, USA, 2011, pp. 69–77 [Online]. Available: <http://doi.acm.org/10.1145/2020408.2020426>
- [35] C. A. Andersson and R. Bro, The N-way Toolbox for Matlab. [Online]. Available: <http://www.models.life.ku.dk/source/nwaytoolbox>
- [36] L. Sorber, M. Van Barel, and L. De Lathauwer, Tensorlab. ver. v2.0 [Online]. Available: <http://www.tensorlab.net/>



Athanasios P. Liavas (M'89) received the Diploma and the Ph.D. from the Department of Computer Engineering and Informatics, University of Patras, Greece, and in 1989 and 1993, respectively. He served as Assistant Professor at the Department of Mathematics of the University of the Aegean (2001–2004); Associate Professor (2004–2009) and Professor (2009–present) at the Department of Electronic and Computer Engineering at the Technical University of Crete. His current research interests lie in the area of signal processing and

machine learning.

Dr. Liavas is a member of the Technical Chamber of Greece.



Nicholas D. Sidiropoulos (F'09) received the Diploma in electrical engineering from the Aristotelian University of Thessaloniki, Greece, and M.S. and Ph.D. degrees in electrical engineering from the University of Maryland—College Park, in 1988, 1990 and 1992, respectively. He served as Assistant Professor at the University of Virginia (1997–1999); Associate Professor at the University of Minnesota—Minneapolis (2000–2002); Professor at the Technical University of Crete, Greece (2002–2011); and Professor at the University of

Minnesota—Minneapolis (2011–), where he currently holds an ADC Chair in Digital Technology. His current research focuses primarily on signal and tensor analytics, with applications in cognitive radio, big data, and preference measurement.

He received the NSF/CAREER award (1998), the IEEE Signal Processing Society (SPS) Best Paper Award (2001, 2007, 2011), and the IEEE SPS Meritorious Service Award (2010). He has served as IEEE SPS Distinguished Lecturer (2008–2009), and Chair of the IEEE Signal Processing for Communications and Networking Technical Committee (2007–2008). He received the Distinguished Alumni Award of the Department of Electrical and Computer Engineering, University of Maryland, College Park (2013). He was elected Fellow of EURASIP in 2014.