

An AO-ADMM approach to constraining PARAFAC2 on all modes*

Marie Roald[†] Carla Schenker[†] Vince D. Calhoun[‡]
Tülay Adalı[§] Rasmus Bro[¶] Jeremy E. Cohen^{||} Evrim Acar^{**}

Abstract

Analyzing multi-way measurements with variations across one mode of the dataset is a challenge in various fields including data mining, neuroscience and chemometrics. For example, measurements may evolve over time or have unaligned time profiles. The PARAFAC2 model has been successfully used to analyze such data by allowing the underlying factor matrices in one mode (i.e., the evolving mode) to change across slices. The traditional approach to fit a PARAFAC2 model is to use an alternating least squares-based algorithm, which handles the constant cross-product constraint of the PARAFAC2 model by implicitly estimating the evolving factor matrices. This approach makes imposing regularization on these factor matrices challenging. There is currently no algorithm to flexibly impose such regularization with general penalty functions and hard constraints. In order to address this challenge and to avoid the implicit estimation, in this paper, we propose an algorithm for fitting PARAFAC2 based on alternating optimization with the alternating direction method of multipliers (AO-ADMM). With numerical experiments on simulated data, we show that the proposed PARAFAC2 AO-ADMM approach allows for flexible constraints, recovers the underlying patterns accurately, and is computationally efficient compared to the state-of-the-art. We also apply our model to two real-world datasets from neuroscience and chemometrics, and show that constraining the evolving mode improves the interpretability of the extracted patterns.

1 Introduction

For many applications in different domains, measurements are obtained in the form of sequences of matrices, which can be arranged as a third-order tensor. Tensor decomposition, which is the higher-order extension of matrix decomposition, is a well-known and useful tool for analysis of such multi-way data. A popular tensor

***Funding:** This work was supported in part by the Research Council of Norway through project 300489 (IKTPLUS) and ANR JCJC project LoRAiA ANR-20-CE23-0010.

[†]Department of Data Science and Knowledge Discovery, Simula Metropolitan Center for Digital Engineering, Oslo, Norway & Faculty of Technology, Art and Design, Oslo Metropolitan University, Oslo, Norway (mariero@simula.no, carla@simula.no).

[‡]Department of Psychology, Georgia State University, Atlanta, GA, USA (vcalhoun@gsu.edu)

[§]Department of Computer Science and Electrical Engineering, UMBC, Baltimore, MD (adali@umbc.edu).

[¶]Department of Food Science, University of Copenhagen, Copenhagen, Denmark (rb@food.ku.dk).

^{||}Univ Lyon, INSA-Lyon, UCBL, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1206, F-69100 Villeurbanne, France (jeremy.cohen@cnrs.fr).

^{**}Department of Data Science and Knowledge Discovery, Simula Metropolitan Center for Digital Engineering, Oslo, Norway (evrim@simula.no).




Figure 1: Illustration of a two-component PARAFAC2 model for a ragged tensor.

decomposition model is the CANDECOMP/PARAFAC (CP) decomposition [2], [3], also known as the canonical polyadic decomposition [4] which describes a tensor as the sum of a minimum number of rank-one components, i.e., an R -component CP model of the tensor $\mathcal{X} \in \mathbb{R}^{I \times J \times K}$, is given by

$$\mathcal{X} \approx \sum_{r=1}^R \mathbf{a}_r \circ \mathbf{b}_r \circ \mathbf{c}_r, \quad (1)$$

where \mathbf{a}_r , \mathbf{b}_r and \mathbf{c}_r are the r -th column of the factor matrices $\mathbf{A} \in \mathbb{R}^{I \times R}$, $\mathbf{B} \in \mathbb{R}^{J \times R}$ and $\mathbf{C} \in \mathbb{R}^{K \times R}$. These factor matrices can be interpreted as a collection of underlying patterns. Following (1), each frontal slice, \mathbf{X}_k of \mathcal{X} can be represented as:

$$\mathbf{X}_k \approx \mathbf{A} \mathbf{D}_k \mathbf{B}^T, \quad (2)$$

where each $\mathbf{D}_k \in \mathbb{R}^{R \times R}$ is diagonal with diagonal entries given by the k -th row of \mathbf{C} . The CP model has successfully extracted meaningful patterns from data tensors in various fields including chemometrics [5], neuroscience [6]–[8] and social network analysis [9] (see surveys on tensor factorizations for more applications [10]–[12]).

However, CP cannot describe patterns that vary across one mode without following the strict assumption in (2). For such evolving patterns, the more general PARAFAC2 model [13] is better suited as it models the slices of a tensor by

$$\mathbf{X}_k \approx \mathbf{A} \mathbf{D}_k \mathbf{B}_k^T, \quad (3)$$

where $\mathbf{D}_k \in \mathbb{R}^{R \times R}$ is diagonal, $\mathbf{A} \in \mathbb{R}^{I \times R}$, $\mathbf{B}_k \in \mathbb{R}^{J \times R}$ and $\mathbf{B}_{k_1}^T \mathbf{B}_{k_1} = \mathbf{B}_{k_2}^T \mathbf{B}_{k_2}$ for all $k_1, k_2 \leq K$. Thus, PARAFAC2 allows the \mathbf{B}_k factor matrices to vary across tensor slices, under the constraint that their cross product is constant, which is more general than CP, where $\mathbf{B}_{k_1} = \mathbf{B}_{k_2}$. Letting \mathbf{B}_k vary in this way also enables decomposing ragged tensors (i.e., stacks of matrices with varying size, as illustrated in Figure 1).

Letting factor matrices evolve across one mode has proven advantageous for applications in many domains. PARAFAC2 has, for example, shown an exceptional ability to analyze chromatographic data with unaligned elution profiles [14]. The PARAFAC2 model has also been successfully applied to resolve unaligned temporal profiles in electronic health records [15] and to retrieve information across languages from a multi-language corpus [16]. In neuroscience, PARAFAC2 has been used to model functional connectivity in functional Magnetic Resonance Imaging (fMRI) data [17], for tracing time-evolving networks of brain connectivity [18] and for jointly analyzing fMRI data from multiple tasks [19].

Online (also referred to as incremental or adaptive) tensor methods can also be used to capture evolving components, e.g., through updates of the CP model [20], [21], in particular, in the context of streaming data. While related, modelling assumptions of online tensor methods vs. the PARAFAC2 model are different. Online methods update the whole model, i.e., factor matrices in all modes, as new data slices arrive at new time points, for instance through efficient stochastic gradient descent (SGD)-based approaches [22]–[24]. On the other hand, PARAFAC2 jointly analyzes data slices summarizing them all using a common \mathbf{A} factor matrix (e.g., when a *subjects* by *voxels* by *time windows* fMRI tensor is analyzed using PARAFAC2, the same subject coefficients, \mathbf{A} , are assumed for all time windows [18]) that can only change up to a scaling across different slices while letting \mathbf{B} factor matrices change from one slice to another. The distinction between modelling assumptions becomes even more clear when multi-task fMRI data (in the form of a *subjects* by *voxels* by *tasks* tensor) is analyzed. PARAFAC2 assumes the same subject coefficients and different voxel-mode (evolving) components for each task [19] whereas online tensor methods are not suitable for the analysis of such data.

One challenge with PARAFAC2 is that the traditional alternating least squares (ALS) approach for fitting the model handles its constant cross-product constraint by implicitly estimating the evolving \mathbf{B}_k factor matrices via the reparametrization $\mathbf{B}_k = \mathbf{P}_k \mathbf{\Delta}_B$ with $\mathbf{P}_k^T \mathbf{P}_k = \mathbf{I}$ [25]. This reparametrization makes it difficult to impose constraints on the evolving \mathbf{B}_k -matrices in a flexible way. Nevertheless, constraints are essential to obtain uniqueness and interpretability in matrix factorizations [26]. Additionally, constraints and regularization can improve the interpretability of components obtained from CP models [27], [28], and the non-evolving modes of PARAFAC2 models [15]. Recent studies have also demonstrated the benefits of constraining the evolving mode of PARAFAC2 [15], [29]–[31]. One way to constrain the \mathbf{B}_k -matrices is by specifying a linear subspace their columns should be contained in. In [29], Helwig showed that the data can be preprocessed by projecting it onto the subspace of interest, before fitting the PARAFAC2 model and that this approach can constrain the evolving components to be spanned by a B-spline basis, thus also constraining them to be smooth [29]. However, a downside of this scheme is that it only supports linear constraints and that we need to know the linear subspace (e.g., the spline knots) of the components a priori, which may be difficult in practice. Another approach to have constraints on \mathbf{B}_k -matrices, in particular, non-negativity constraints, is the flexible coupling approach [30]. This approach fits a non-negative coupled matrix factorization model with a regularization term based on the PARAFAC2 constraint. While this approach, in theory, can be used with any constrained least squares solvers, the algorithm in [30] uses hierarchical alternating least squares (HALS) and supports only non-negativity constraints. Other notable approaches include LogPar [32] using another regularization penalty based on PARAFAC2 to improve uniqueness properties of regularized coupled non-negative matrix factorization for binary data, and CANDELINC2 [31], which uses a scheme combining CANDELINC (Canonical Decomposition with Linear Constraints) [33]–[35] and non-negativity constraints for PARAFAC2. Despite these recent efforts, available methods for PARAFAC2 are still limited in terms of the type of constraints they can impose.

In this paper, we propose fitting PARAFAC2 using an alternating optimization (AO) scheme with the alternating direction method of multipliers (ADMM) to facilitate the use of a wider set of constraints in all modes (both evolving and non-evolving). Recently, Huang et al. introduced the AO-ADMM scheme for constrained CP models [36], and Schenker et al. extended this framework to regularized linearly coupled matrix-tensor decompositions [37], [38]. AO-ADMM has been successfully used to im-

pose proximable constraints on the non-evolving factor matrices of the PARAFAC2 model [15], [39]. However, using that approach, evolving factor matrices are still modelled implicitly by $\mathbf{B}_k = \mathbf{P}_k \mathbf{\Delta}_B$ as in [25], with AO-ADMM being used for the CP step of the algorithm. Thus, only linear constraints can be imposed directly on the evolving mode, still limiting the possible constraints on the evolving factor matrices. Unlike earlier studies, we present an AO-ADMM based framework for the PARAFAC2 model with ADMM updates that widen the possible constraints to any proximable penalty function in all modes. With numerical experiments on both simulated and real data, we demonstrate the performance of this framework in terms of

- **Flexibility:** The scheme allows for regularization of the evolving components with any proximable penalty function.
- **Efficiency:** The scheme can fit non-negative PARAFAC2 models faster than the flexible coupling with HALS approach introduced in [30].
- **Accuracy:** Applying suitable constraints for the evolving components can improve performance compared to only constraining the non-evolving modes.

This paper extends our preliminary study [1], where we introduced and showed the promise of AO-ADMM for constraining PARAFAC2 models using non-negativity constraints, graph Laplacian and total variation regularization. In this paper, we provide a more detailed derivation and a theoretical discussion of both the AO-ADMM scheme and the regularized PARAFAC2 problem. We include extensive numerical experiments using more general simulation setups and one new constraint (unimodal component vectors). Finally, we demonstrate the usefulness of the proposed algorithmic approach on two applications from neuroimaging and chemometrics.

Section 2 briefly states the notation used in this paper, and Section 3 introduces PARAFAC2 and states some observations on the PARAFAC2 constraint and available methods to fit PARAFAC2 models. In Section 4, we introduce the new AO-ADMM scheme that enables regularization on all modes. Numerical experiments evaluating the proposed scheme on seven simulation setups and two real datasets are presented in Section 5. Finally, Section 6 discusses our results and possible future work.

2 Notation

We specify some key tensor concepts and notations here and refer to [10] for a thorough review of the nomenclature and background theory. A tensor can be considered a multidimensional array that extends the concept of matrices to higher dimensions. The dimensions of a tensor are called modes, and the number of modes is referred to as the order. Thus, vectors are first-order tensors, matrices are second-order tensors, “cubes” of numbers are third-order tensors and so forth. Tensors with three modes or more are often called *higher-order tensors*. We indicate vectors by bold lowercase letters, e.g., \mathbf{v} , matrices by bold uppercase letters, e.g., \mathbf{M} , and higher-order tensors by bold uppercase calligraphic letters, e.g., \mathcal{X} . Furthermore, we denote the r -th column of a matrix, \mathbf{X} , by \mathbf{x}_r and the r -th column of a matrix \mathbf{X}_k by $[\mathbf{x}_k]_r$. Sets of matrices are denoted by calligraphic letters (e.g., $\mathcal{O}_{n,m} = \{\mathbf{P} \in \mathbb{R}^{n \times m} : \mathbf{P}^\top \mathbf{P} = \mathbf{I}\}$) and sets of sets of matrices are denoted by script letters (e.g., \mathcal{P}).

All tensors we consider in this work have three modes or less¹. The Khatri-Rao

¹But our work extends to any order trivially. For orders above three, the PARAFAC2 constraint is not needed for uniqueness [14] since a fourth-order PARAFAC2 model is equivalent to a coupled CP model. Still, the PARAFAC2 constraint can be used to enforce a constant relationship between the components for each frontal slab (higher-order slice) of the data tensor.

product (column-wise Kronecker product) is denoted by \odot , the Hadamard product (element-wise multiplication) is denoted by $*$, the vector outer product is denoted by \circ and the Frobenius norm is denoted by $\|\cdot\|_F$. The matrix inverse and transpose are denoted by $^{-1}$ and $^\top$, respectively. The number of components for tensor decomposition models is denoted by R , and I, J and K denote the number of horizontal, lateral and frontal slices in a tensor. For ragged tensors (i.e., stacks of matrices with varying number of columns), J_k represents the number of columns for the k -th slice.

3 PARAFAC2 & ALS

Recall that the PARAFAC2 model represents \mathbf{X}_k as:

$$\mathbf{X}_k \approx \mathbf{A} \mathbf{D}_k \mathbf{B}_k^\top, \quad (4)$$

where the $\mathbf{B}_k \in \mathbb{R}^{J \times R}$ matrices follow the *PARAFAC2 constraint*. Specifically, $\mathbf{B}_{k_1}^\top \mathbf{B}_{k_1} = \mathbf{B}_{k_2}^\top \mathbf{B}_{k_2}$ for all $k_1, k_2 \leq K$ (henceforth, $\{\mathbf{B}_k\}_{k \leq K} \in \mathcal{P}$). If one considers the columns of the factor matrices $\{\mathbf{B}_k\}_{k \leq K}$ as some underlying patterns or sources underlying measurements, then PARAFAC2 allows these patterns to rotate and mirror for each slice k while the cross product between these patterns stays constant. Moreover, the constant cross-product constraint is equivalent to a constant correlation constraint if the columns of the \mathbf{B}_k matrices sum to zero (which can be ensured using the method in [29]). PARAFAC2 has also been shown to extract meaningful patterns, even when the data does not fully follow the PARAFAC2 constraint [14], [18]. The next section states some further mathematical properties of this constraint.

3.1 Observations on the PARAFAC2 constraint

The set

$$\mathcal{P} = \{ \{\mathbf{B}_k\}_{k=1}^K \mid \mathbf{B}_{k_1}^\top \mathbf{B}_{k_1} = \mathbf{B}_{k_2}^\top \mathbf{B}_{k_2} \ \forall k_1, k_2 \leq K \} \quad (5)$$

defines the PARAFAC2 constraint, and restricts the angle between component vectors of \mathbf{B}_k to be constant for all k . Without this constraint, the PARAFAC2 model would be equivalent to the rank R matrix factorization of the stacked matrix $[\mathbf{X}_1, \dots, \mathbf{X}_K]$ whose parameters are not identifiable in general. The PARAFAC2 constraint ensures unique components given enough frontal slices, \mathbf{X}_k [25], [40], [41]. Despite its importance, the set \mathcal{P} and its properties have received little attention in the PARAFAC2 literature.

Below, we make a few key observations on the PARAFAC2 constraint. First, as shown by Kiers et al.[25], given a collection of matrices $\{\mathbf{B}_k\}_{k=1}^K \in \mathcal{P}$, there exists a matrix $\mathbf{\Delta}_B \in \mathbb{R}^{R \times R}$ and a collection $\{\mathbf{P}_k\}_{k=1}^K$ of orthogonal matrices such that

$$\forall k \leq K, \mathbf{B}_k = \mathbf{P}_k \mathbf{\Delta}_B. \quad (6)$$

This means that \mathcal{P} is related to $\mathcal{O}_{J_1, R} \times \dots \times \mathcal{O}_{J_K, R} \times \mathbb{R}^{R \times R}$. We use this fact in the supplementary material to prove that \mathcal{P} is closed.

Proposition 3.1. *The set \mathcal{P} is closed.*

As a direct corollary, the projection on \mathcal{P} is well-defined and generally unique [42]. Projecting on \mathcal{P} will be of importance in Section 4.

Second, it may also be noted that \mathcal{P} is the set of all discretizations with K points of elements in $\Omega = \{ \mathcal{O}_{R, R} \mathbf{\Delta}_B \mid \mathbf{\Delta}_B \in \mathbb{R}^{R \times R} \}^2$, where $\mathcal{O}_{R, R} \mathbf{\Delta}_B = \{ \mathbf{B}_k \in \mathbb{R}^{R \times R} \mid \exists \mathbf{P}_k \in$

²One may assume *w.l.o.g.* that $J_k = R$, up to dimensionality reduction of each slice \mathbf{X}_k

$\mathcal{O}_{R,R}, \mathbf{B}_k = \mathbf{P}_k \Delta_{\mathbf{B}}\}$. Studying this continuous set Ω rather than the discrete \mathcal{P} provides another interesting point of view. Indeed, projecting on the PARAFAC2 constraint also means finding the set $\mathcal{O}_{R,R} \Delta_{\mathbf{B}}$ which is the closest to all matrices \mathbf{B}_k . In fact, noticing that the point-set distance satisfies

$$d(\mathbf{B}_k, \mathcal{O}_{R,R} \Delta_{\mathbf{B}}) = \min_{\mathbf{P}_k \in \mathcal{O}_{R,R}} \|\mathbf{B}_k - \mathbf{P}_k \Delta_{\mathbf{B}}\|_F^2 = d(\Delta_{\mathbf{B}}, \mathcal{O}_{R,R} \mathbf{B}_k), \quad (7)$$

one may rewrite the projection problem on \mathcal{P} as the following minimization problem:

$$\min_{\Delta_{\mathbf{B}}} \sum_{k=1}^K d(\Delta_{\mathbf{B}}, \mathcal{O}_{R,R} \mathbf{B}_k). \quad (8)$$

Intuitively, projecting on the PARAFAC2 constraint therefore amounts to finding the closest point $\Delta_{\mathbf{B}}$ to a set of curves $\mathcal{O}_{R,R} \mathbf{B}_k$. Because the projection on \mathcal{P} is in fact a bilevel optimization problem, this justifies the empirical use of an alternating algorithm to compute the projection on \mathcal{P} which is detailed in Section 4. We also hope this geometric point of view can lead to a better understanding of the PARAFAC2 constraint. A striking example is the rank one case $R = 1$, where Ω simply becomes the set of all spheres, and projecting on \mathcal{P} amounts to projecting each \mathbf{B}_k on the average sphere of radius $\frac{1}{K} \sum_{k=1}^K \|\mathbf{B}_k\|_2$. Some illustrations and the computations for the rank one case are provided in the supplementary materials.

3.2 Recoverability of the evolving components

While the PARAFAC2 model has many similarities with the CP model, it comes with some extra challenges. One notable challenge is the additional sign indeterminacy of the PARAFAC2 model. Specifically, for any component, r , we can reverse the sign of $[d_k]_{rr}$ together with $[\mathbf{b}_k]_r$, so long as the corresponding signs for all components, s , that satisfy $[\mathbf{b}_k]_r^T [\mathbf{b}_k]_s \neq 0$ are also reversed [13], [43]. There are several ways to handle this sign indeterminacy, e.g. by imposing non-negativity constraints on $\{\mathbf{D}_k\}_{k \leq K}$ [13], aligning the components with the direction of the data [43] or by using additional information [44].

Additionally, for real data matrices that are affected by noise, the recoverability of the r -th column of \mathbf{B}_k is directly affected by the magnitude of the r -th diagonal entry in \mathbf{D}_k . To see this, consider estimating \mathbf{B}_k from noisy data $\tilde{\mathbf{X}}_k = \mathbf{X}_k + \mathbf{E}_k$ with $\mathbf{X}_k = \mathbf{A} \mathbf{D}_k \mathbf{B}_k^T$ and known \mathbf{A} and \mathbf{D}_k . Using the normal equation for \mathbf{B}_k , we get

$$\hat{\mathbf{B}}_k^T = \mathbf{D}_k^{-1} (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T (\mathbf{X}_k + \mathbf{E}_k) = \mathbf{B}_k^T + \mathbf{D}_k^{-1} (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{E}_k. \quad (9)$$

This means that the estimate of the r -th column of \mathbf{B}_k , $[\hat{\mathbf{b}}_k]_r$, is affected by the noise scaled by $1/[d_k]_{rr}$ (for simplicity, we omit $(\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}$, since it is constant for all k). Thus, if an element of \mathbf{D}_k is small compared to the noise strength $\|\mathbf{E}_k\|_F$, then the corresponding column of \mathbf{B}_k will be poorly estimated.

To quantify the effect of the noise on recovering a column of \mathbf{B}_k , we define the *column-wise signal-to-noise ratio* (cwSNR), defined by

$$\text{cwSNR}_{kr} = \frac{[d_k]_{rr}^2}{\|\mathbf{E}_k\|^2}, \quad (10)$$

with the factor matrices scaled so \mathbf{A} and \mathbf{B}_k have unit norm columns. We show that the cwSNR is a good predictor on the accuracy of \mathbf{B}_k estimates in Section 5.2.6.

3.3 ALS for PARAFAC2 (unregularized)

There is no closed form solution for fitting PARAFAC2 models. Instead, one has to solve a difficult, non-convex, optimization problem. This is usually done by estimating the solution using *alternating optimization* (AO), which, instead of solving the optimization problem directly, fixes all but one factor matrix in each step. This single factor matrix can then be updated efficiently by solving a quadratic surrogate problem.

We cannot easily use AO directly on Equation (4), as the PARAFAC2 constraint leads to a non-convex optimization problem for the \mathbf{B}_k updates. However, as discussed in Section 3.1, the PARAFAC2 model can be reformulated as:

$$\mathbf{X}_k \approx \mathbf{A} \mathbf{D}_k \mathbf{\Delta}_B^\top \mathbf{P}_k^\top, \quad (11)$$

where $\mathbf{\Delta}_B \in \mathbb{R}^{R \times R}$ and $\mathbf{P}_k \in \mathcal{O}_{J_k, R}$. This problem can be solved efficiently using an alternating least squares procedure. In this case, the \mathbf{P}_k updates are performed by solving an orthogonal Procrustes problem, resulting in Algorithm 1 [25].

Algorithm 1: PARAFAC2 ALS [25]

Result: \mathbf{A} , $\{\mathbf{P}_k\}_{k \leq K}$, $\mathbf{\Delta}_B$, $\{\mathbf{D}_k\}_{k \leq K}$

- 1 Initialize \mathbf{A} , $\mathbf{\Delta}_B$ and $\{\mathbf{D}_k\}_{k \leq K}$
 - 2 **while** *stopping conditions are not met and max no. iterations not exceeded*
 - 3 **do**
 - 4 **for** $k \leftarrow 1$ **to** K **do**
 - 5 Compute a rank R truncated SVD: $\mathbf{U}_k \mathbf{S}_k \mathbf{V}_k^\top = \mathbf{X}_k^\top \mathbf{A} \mathbf{D}_k \mathbf{\Delta}_B^\top$
 - 6 $\mathbf{P}_k \leftarrow \mathbf{U}_k \mathbf{V}_k^\top$
 - 7 $\mathbf{T}_k \leftarrow \mathbf{X}_k \mathbf{P}_k$
 - 8 **end**
 - 9 Estimate \mathbf{A} , $\mathbf{\Delta}_B$ and $\{\mathbf{D}_k\}_{k \leq K}$ by running a few iterations (e.g., 5) of ALS to fit an R -component CP model to the tensor \mathcal{J} with frontal slices given by \mathbf{T}_k .
 - 10 **end**
-

3.4 Constrained PARAFAC2 with flexible coupling

The ALS formulation of PARAFAC2 does not lend itself for constraining or regularizing the \mathbf{B}_k components. Currently the only way to impose non-negativity on these components is with the *flexible coupling approach* by Cohen and Bro [30]. This approach works by fitting a coupled matrix factorization with a regularization term that penalizes the distance between $\{\mathbf{B}_k\}_{k \leq K}$ and a point on a curve $\mathcal{O}_{R, R} \mathbf{\Delta}_B$ whose coordinates are also parameters to estimate, yielding the following optimization problem:

$$\begin{aligned} & \min_{\mathbf{A}, \{\mathbf{D}_k\}_{k \leq K}, \{\mathbf{B}_k\}_{k \leq K}, \mathbf{\Delta}_B, \{\mathbf{P}_k\}_{k \leq K}} \left[\sum_{k=1}^K \left\| \mathbf{X}_k - \mathbf{A} \mathbf{D}_k \mathbf{B}_k^\top \right\|_F^2 + \mu_t \left\| \mathbf{B}_k - \mathbf{P}_k \mathbf{\Delta}_B \right\|_F^2 \right] \\ & \text{subject to } \mathbf{P}_k^\top \mathbf{P}_k = \mathbf{I} \quad \forall k \leq K, \\ & \mathbf{A} \geq 0, \\ & \mathbf{B}_k, \mathbf{D}_k \geq 0 \quad \forall k \leq K, \end{aligned} \quad (12)$$

which can be solved with HALS [45], also known as column-wise updates [46], [47]. To ensure that the components satisfy the PARAFAC2 constraint, the regularization parameter, μ_t , is adaptively selected and increased after every iteration. The word “flexible” in the name of this approach stems from the fact that the \mathbf{B}_k matrices are not directly parametrized by $\mathbf{P}_k \mathbf{\Delta}_B$, but instead encouraged to be close to this product through a “flexible coupling”. While the trick of gradually increasing the regularization strength to obtain models that follow the PARAFAC2 constraint could, in theory, be used with any constrained least squares solver, the algorithm in [30] uses HALS and supports only non-negativity constraints. Furthermore, to the best of our knowledge, this formulation has only ever been used with non-negativity constraints.

4 PARAFAC2 AO-ADMM

Consider the regularized PARAFAC2 problem:

$$\begin{aligned} \min_{\mathbf{A}, \{\mathbf{B}_k, \mathbf{D}_k\}_{k \leq K}} & \left[f(\mathbf{A}, \{\mathbf{B}_k\}_{k \leq K}, \{\mathbf{D}_k\}_{k \leq K}) + g_{\mathbf{A}}(\mathbf{A}) + \sum_{k=1}^K \{g_{\mathbf{B}_k}(\mathbf{B}_k) + g_{\mathbf{D}_k}(\mathbf{D}_k)\} \right] \\ \text{subject to } & \{\mathbf{B}_k\}_{k \leq K} \in \mathcal{P}, \end{aligned} \tag{13}$$

where f is the sum of squared errors data fidelity function, given by

$$f(\mathbf{A}, \{\mathbf{B}_k\}_{k \leq K}, \{\mathbf{D}_k\}_{k \leq K}) = \sum_{k=1}^K \left\| \mathbf{A} \mathbf{D}_k \mathbf{B}_k^T - \mathbf{X}_k \right\|_F^2, \tag{14}$$

and $g_{\mathbf{A}}$, $g_{\mathbf{B}_k}$, $g_{\mathbf{D}_k}$ are proximable regularization penalties.

As mentioned above, we cannot easily solve Equation (13) using the formulation in Equation (11), since it would require solving a regularized least squares problem with orthogonality constraints. Thus, we instead propose to solve Equation (13) with a scheme based on ADMM, solving each subproblem approximately. Specifically, we use an AO-ADMM algorithm.

AO-ADMM has been used for fitting constrained CP models [36] and coupled matrix and tensor factorizations [37], [38]. It has also been used for fitting PARAFAC2 models, by using AO-ADMM to update \mathbf{A} , $\mathbf{\Delta}_B$ and \mathbf{C} (corresponding to using AO-ADMM instead of ALS to fit the CP model in Line 8 of Algorithm 1) [15]. However, with such a scheme, the \mathbf{B}_k matrices can only be constrained linearly.

We instead propose to use AO-ADMM to solve Equation (13) without reparametrization, thus allowing for proximal regularization of all factor matrices including $\{\mathbf{B}_k\}_{k \leq K}$. For this, we introduce an algorithm for projecting onto \mathcal{P} , which makes it feasible to use ADMM for solving regularized least squares problems with the PARAFAC2 constraint, and obtain the PARAFAC2 AO-ADMM algorithm described in Algorithm 2.

In the following sections, we describe how each step of this algorithm is conducted in detail. First, we summarize ADMM in Section 4.1 and AO-ADMM in Section 4.2. Then, Section 4.3 and Section 4.4, describes ADMM update steps for lines 3 and 4-5 of Algorithm 2, respectively. Section 4.5 outlines stopping conditions and a heuristic to compute appropriate feasibility penalty parameters. Finally, Section 4.6 discusses the computational complexity for each step of the algorithm and Section 4.7 lists some proximable penalties (with more details in Section SM3) and notes some important considerations when using penalty-based regularization.

Algorithm 2: AO-ADMM for PARAFAC2

Result: $\mathbf{A}, \{\mathbf{B}_k, \mathbf{D}_k\}_{k \leq K}$

- 1 Initialize $\mathbf{A}, \mathbf{Z}_A, \boldsymbol{\mu}_A, \mathbf{B}_k, \mathbf{Z}_{B_k}, \boldsymbol{\mu}_{Z_{B_k}}, \Delta_B, \mathbf{P}_k, \boldsymbol{\mu}_{\Delta_{B_k}}, \mathbf{D}_k, \mathbf{Z}_{D_k}$, and $\boldsymbol{\mu}_{D_k}$
 - 2 **while** *stopping conditions are not met and max no. iterations not exceeded*
do
 - 3 Update $\left\{ \mathbf{B}_k, \mathbf{Z}_{B_k}, \mathbf{P}_k, \boldsymbol{\mu}_{Z_{B_k}}, \boldsymbol{\mu}_{\Delta_{B_k}} \right\}_{k \leq K}$ and Δ_B using Algorithm 4
 - 4 Update \mathbf{A}, \mathbf{Z}_A and $\boldsymbol{\mu}_A$ using Algorithm 6
 - 5 Update $\left\{ \mathbf{D}_k, \mathbf{Z}_{D_k}, \boldsymbol{\mu}_{D_k} \right\}_{k \leq K}$ using Algorithm 7
 - 6 **end**
-

4.1 Summary of ADMM

ADMM solves optimization problems in the form:

$$\begin{aligned} & \min_{\mathbf{x}, \mathbf{y}} [f(\mathbf{x}) + g(\mathbf{y})] \\ & \text{subject to } \mathbf{M}\mathbf{x} + \mathbf{N}\mathbf{y} = \mathbf{c}, \end{aligned} \quad (15)$$

where \mathbf{M} and \mathbf{N} are known matrices and \mathbf{c} is a known vector. For ADMM, we first formulate the augmented Lagrange dual problem of Equation (15):

$$\max_{\boldsymbol{\nu}} \left[\min_{\mathbf{x}, \mathbf{y}} \left[f(\mathbf{x}) + g(\mathbf{y}) + \boldsymbol{\nu}^\top (\mathbf{M}\mathbf{x} + \mathbf{N}\mathbf{y} - \mathbf{c}) + \frac{\rho}{2} \|\mathbf{M}\mathbf{x} + \mathbf{N}\mathbf{y} - \mathbf{c}\|_2^2 \right] \right], \quad (16)$$

where ρ is a penalty parameter that defines the penalty for infeasible solutions. Then, we solve the saddle-point problem Equation (16) with AO on \mathbf{x} and \mathbf{y} and gradient ascent on $\boldsymbol{\nu}$. Defining the scaled dual-variable, $\boldsymbol{\mu} = \boldsymbol{\nu}/\rho$, we obtain Algorithm 3.

One benefit is that whenever \mathbf{M} or \mathbf{N} are identity matrices, then the corresponding update step reduces to evaluating the scaled proximal operator (shown for g):

$$\text{prox}_{\frac{g}{\rho}}(-(\mathbf{M}\mathbf{x} - \mathbf{c} + \boldsymbol{\mu})) = \arg \min_{\mathbf{y}} \left[g(\mathbf{y}) + \frac{\rho}{2} \|\mathbf{y} + \mathbf{M}\mathbf{x} - \mathbf{c} + \boldsymbol{\mu}\|_2^2 \right], \quad (17)$$

which can be evaluated efficiently for a large family of functions. The update steps for g_A , g_{B_k} , and g_{D_k} in Equation (13) can always be reduced to this form, enabling any proximal regularization on all modes. See Section 4.7 for details on proximal operators and examples of proximal penalties. For a thorough review of ADMM, we refer to [48].

Algorithm 3: ADMM [48]

Result: $\mathbf{x}, \mathbf{y}, \boldsymbol{\mu}$

- 1 Initialize \mathbf{x}, \mathbf{y} and $\boldsymbol{\mu}$
 - 2 **while** *stopping conditions are not met and max no. iterations not exceeded*
do
 - 3 $\mathbf{x} \leftarrow \arg \min_{\mathbf{x}} \left[f(\mathbf{x}) + \frac{\rho}{2} \|\mathbf{M}\mathbf{x} + \mathbf{N}\mathbf{y} - \mathbf{c} + \boldsymbol{\mu}\|_2^2 \right]$
 - 4 $\mathbf{y} \leftarrow \arg \min_{\mathbf{y}} \left[g(\mathbf{y}) + \frac{\rho}{2} \|\mathbf{M}\mathbf{x} + \mathbf{N}\mathbf{y} - \mathbf{c} + \boldsymbol{\mu}\|_2^2 \right]$
 - 5 $\boldsymbol{\mu} \leftarrow \mathbf{M}\mathbf{x} + \mathbf{N}\mathbf{y} - \mathbf{c} + \boldsymbol{\mu}$
 - 6 **end**
-

4.2 Overview of AO-ADMM

Consider now a problem in the form

$$\min_{\mathbf{x}, \mathbf{y}} [f(\mathbf{x}, \mathbf{y}) + g_x(\mathbf{x}) + g_y(\mathbf{y})], \quad (18)$$

where the proximal operator for f is costly to evaluate, but the proximal operators of $f_{\mathbf{x}}(\mathbf{y}) = f(\mathbf{x}, \mathbf{y})$, $f_{\mathbf{y}}(\mathbf{x}) = f(\mathbf{x}, \mathbf{y})$, $g_{\mathbf{x}}$ and $g_{\mathbf{y}}$ are easy to evaluate. In that case, we can use an AO-scheme, where we fix \mathbf{y} and update \mathbf{x} with several iterations of ADMM, and vice versa. Such a scheme is called an AO-ADMM scheme [36].

There are, to the best of our knowledge, no convergence guarantees for AO-ADMM. For AO, it is known that if $f_{\mathbf{y}}(\mathbf{x}) + g_x(\mathbf{x})$ and $f_{\mathbf{x}}(\mathbf{y}) + g_y(\mathbf{y})$ always have unique minima and these minima are attained for each iteration, then all limit points of the algorithm are stationary points (see Proposition 2.7.1 in [49] or the discussion of convergence in [36]). Thus, if we impose a strictly convex regularization penalty (e.g., a ridge penalty), the AO-ADMM algorithm would provably converge to a stationary point given infinitely many ADMM iterations for each block update. While we often have convex regularization for updating the \mathbf{A} - and \mathbf{C} -matrix, the PARAFAC2 constraint for the \mathbf{B}_k -matrices is not convex, which makes it hard to conclude about the convergence of the ADMM updates for $\{\mathbf{B}_k\}_{k \leq K}$, and therefore also the AO-ADMM algorithm. Still, in our experience, the AO-ADMM algorithm converges in most cases.

4.3 An ADMM scheme for $\{\mathbf{B}_k\}_{k \leq K}$ in Algorithm 2

Our main contribution is introducing an ADMM scheme to solve the non-convex optimization problem

$$\begin{aligned} \min_{\{\mathbf{B}_k\}_{k \leq K}} & \left[\sum_{k=1}^K f_{\mathbf{B}_k}(\mathbf{B}_k) + g_{\mathbf{B}_k}(\mathbf{B}_k) \right] \\ \text{subject to} & \quad \{\mathbf{B}_k\}_{k \leq K} \in \mathcal{P}, \end{aligned} \quad (19)$$

where $f_{\mathbf{B}_k}(\mathbf{B}_k) = \|\mathbf{A}\mathbf{D}_k\mathbf{B}_k^\top - \mathbf{X}_k\|_F^2$. We introduce the following splitting scheme

$$\begin{aligned} \min_{\{\mathbf{B}_k\}_{k \leq K}} & \left[\sum_{k=1}^K \{f_{\mathbf{B}_k}(\mathbf{B}_k) + g_{\mathbf{B}_k}(\mathbf{Z}_{\mathbf{B}_k})\} + \iota_{\mathcal{P}}(\{\mathbf{Y}_{\mathbf{B}_k}\}_{k \leq K}) \right] \\ \text{subject to} & \quad \mathbf{B}_k = \mathbf{Z}_{\mathbf{B}_k} \quad \forall k \leq K, \\ & \quad \mathbf{B}_k = \mathbf{Y}_{\mathbf{B}_k} \quad \forall k \leq K, \end{aligned} \quad (20)$$

where $\iota_{\mathcal{P}}(\{\mathbf{Y}_{\mathbf{B}_k}\}_{k \leq K}) = 0$ if $\{\mathbf{Y}_{\mathbf{B}_k}\}_{k \leq K} \in \mathcal{P}$, and ∞ otherwise. This splitting scheme is in the standard form for problems that are solvable with ADMM, which we can utilize to obtain Algorithm 4. Next, we describe each line in this algorithm in detail.

Two main lines in Algorithm 4 need to be derived: the update rules for $\{\mathbf{B}_k\}_{k \leq K}$ in Line 3 and the update rules for $\{\mathbf{Y}_{\mathbf{B}_k}\}_{k \leq K}$ in Line 6. The proximal operator for $g_{\mathbf{B}_k}$ in Line 4 is specified by the regularization function. The update rule for \mathbf{B}_k requires us to solve a least squares problem, which has a closed form solution:

$$\mathbf{B}_k \leftarrow \left(\mathbf{X}_k^\top \mathbf{A} \mathbf{D}_k + \frac{\rho_{\mathbf{B}_k}}{2} \left(\mathbf{Z}_{\mathbf{B}_k} - \boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}} + \mathbf{Y}_{\mathbf{B}_k} - \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}} \right) \right) \left(\mathbf{D}_k \mathbf{A}^\top \mathbf{A} \mathbf{D}_k + \rho_{\mathbf{B}_k} \mathbf{I} \right)^{-1}. \quad (21)$$

while the update rules for $\{\mathbf{Y}_{\mathbf{B}_k}\}_{k \leq K}$ require an efficient way to estimate $\text{prox}_{\iota_{\mathcal{P}}}$.

Algorithm 4: ADMM updates for the \mathbf{B}_k -matrices

Result: $\mathbf{B}_k, \mathbf{Z}_{\mathbf{B}_k}, \mathbf{Y}_{\mathbf{B}_k}, \boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}}, \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}}$

- 1 **while** *stopping conditions are not met and max no. iterations not exceeded*
- do**
- 2 **for** $k \leftarrow 1$ **to** K **do**
- 3 $\mathbf{B}_k \xleftarrow{\text{Equation (21)}} \arg \min_{\mathbf{B}_k} \left[\begin{aligned} & f_{\mathbf{B}_k}(\mathbf{B}_k) + \frac{\rho_{\mathbf{B}_k}}{2} \left\| \mathbf{B}_k - \mathbf{Z}_{\mathbf{B}_k} + \boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}} \right\|^2 \\ & + \frac{\rho_{\mathbf{B}_k}}{2} \left\| \mathbf{B}_k - \mathbf{Y}_{\mathbf{B}_k} + \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}} \right\|^2 \end{aligned} \right]$
- 4 $\mathbf{Z}_{\mathbf{B}_k} \leftarrow \text{prox}_{\frac{\rho_{\mathbf{B}_k}}{\rho_{\mathbf{B}_k}}}(\mathbf{B}_k + \boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}})$
- 5 **end**
- 6 $\{\mathbf{Y}_{\mathbf{B}_k}\}_{k \leq K} \xleftarrow{\text{Alg. 5}} \text{prox}_{\ell, \mathcal{P}}(\{\mathbf{B}_k + \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}}\}_{k \leq K})$
- 7 **for** $k \leftarrow 1$ **to** K **do**
- 8 $\boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}} \leftarrow \mathbf{B}_k - \mathbf{Z}_{\mathbf{B}_k} + \boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}}$
- 9 $\boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}} \leftarrow \mathbf{B}_k - \mathbf{Y}_{\mathbf{B}_k} + \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}}$
- 10 **end**
- 11 **end**

Evaluating $\text{prox}_{\ell, \mathcal{P}}$ is equivalent to evaluating a projection onto \mathcal{P} , which we estimate using the parametrization of \mathcal{P} discussed in Section 3.1, yielding:

$$\begin{aligned} & \min_{\Delta_{\mathbf{B}}, \{\mathbf{P}_k\}_{k \leq K}} \left[\sum_{k=1}^K \frac{\rho_{\mathbf{B}_k}}{2} \left\| \mathbf{B}_k - \mathbf{P}_k \Delta_{\mathbf{B}} + \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}} \right\|^2 \right] \\ & \text{subject to } \mathbf{P}_k^T \mathbf{P}_k = \mathbf{I} \quad \forall k \leq K. \end{aligned} \quad (22)$$

This equation can be approximated efficiently with an AO scheme, where the orthogonal Procrustes problem for each \mathbf{P}_k is solved independently. With this, we obtain Algorithm 5. In our experience, one iteration in Algorithm 5 for each ADMM iteration in Algorithm 4 is sufficient.

Algorithm 5: Approximate projection onto \mathcal{P}

Result: $\{\mathbf{P}_k\}_{k \leq K}, \Delta_{\mathbf{B}}$

- 1 **while** *stopping conditions are not met and max no. iterations not exceeded*
- do**
- 2 **for** $k \leftarrow 1$ **to** K **do**
- 3 Compute ‘‘economy style’’ SVD: $(\mathbf{B}_k + \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}}) \Delta_{\mathbf{B}}^T = \mathbf{U}^{(k)} \boldsymbol{\Sigma}^{(k)} \mathbf{V}^{(k)T}$
- 4 $\mathbf{P}_k \leftarrow \mathbf{U}^{(k)} \mathbf{V}^{(k)T}$
- 5 **end**
- 6 $\Delta_{\mathbf{B}} \leftarrow \frac{1}{\sum_{k=1}^K \rho_{\mathbf{B}_k}} \sum_{k=1}^K \rho_{\mathbf{B}_k} \mathbf{P}_k^T (\mathbf{B}_k + \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}})$
- 7 **end**

4.4 ADMM updates for \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ in Algorithm 2

ADMM updates for \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ can be obtained in multiple ways. Here, we introduce ADMM schemes based on the coupled matrix decomposition interpretation

of PARAFAC2 (CMF-based updates). However, by assuming primal feasibility in Equation (20) (i.e., $\mathbf{B}_k = \mathbf{P}_k \mathbf{\Delta}_k$), we can use CP-based updates for \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ [15]. For discussions and comparisons with the CP-based scheme, see Section SM5.

For the CMF-based update steps, we want to solve the optimization problems

$$\min_{\mathbf{A}} [f_{\mathbf{A}}(\mathbf{A}) + g_{\mathbf{A}}(\mathbf{A})] \quad \text{and} \quad \min_{\mathbf{D}_k} [f_{\mathbf{D}_k}(\mathbf{D}_k) + g_{\mathbf{D}_k}(\mathbf{D}_k)], \quad (23)$$

for updating \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, respectively. $f_{\mathbf{A}}$ and $f_{\mathbf{D}_k}$ are the sum of squared errors data fidelity function for \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, respectively, and $g_{\mathbf{A}}$ and $g_{\mathbf{D}_k}$ are regularization penalties. We apply ADMM directly to these optimization problems.

The proximal operators for the data fidelity functions consist of solving least squares problems. The proximal operator with respect to \mathbf{A} , is given by

$$\text{prox}_{\frac{f_{\mathbf{A}}}{\rho_{\mathbf{A}}}}(\mathbf{M}) = \left(\sum_{k=1}^K \mathbf{X}_k \mathbf{\Gamma}_k + \frac{\rho_{\mathbf{A}}}{2} \mathbf{M} \right) \left(\sum_{k=1}^K \mathbf{\Gamma}_k^{\top} \mathbf{\Gamma}_k + \frac{\rho_{\mathbf{A}}}{2} \mathbf{I} \right)^{-1}, \quad (24)$$

where $\mathbf{\Gamma}_k = \mathbf{B}_k \mathbf{D}_k$. To compute $\text{prox}_{\frac{f_{\mathbf{D}_k}}{\rho_{\mathbf{D}_k}}}(\mathbf{v})$, we consider the vectorized problem:

$$\text{prox}_{\frac{f_{\mathbf{D}_k}}{\rho_{\mathbf{D}_k}}}(\mathbf{v}) = \arg \min_{\mathbf{c}_{k:}} \left[\left\| (\mathbf{A} \odot \mathbf{B}_k) \mathbf{c}_{k:}^{\top} - \text{vec}(\mathbf{X}_k) \right\|_F^2 + \frac{\rho_{\mathbf{D}_k}}{2} \left\| \mathbf{c}_{k:}^{\top} - \mathbf{v} \right\|_F^2 \right], \quad (25)$$

where $\mathbf{c}_{k:}^{\top}$ is a vector containing the diagonal entries of \mathbf{D}_k . Then, we use $(\mathbf{A} \odot \mathbf{B}_k)^{\top} (\mathbf{A} \odot \mathbf{B}_k) = (\mathbf{A}^{\top} \mathbf{A} * \mathbf{B}_k^{\top} \mathbf{B}_k)$ to obtain

$$\text{prox}_{\frac{f_{\mathbf{D}_k}}{\rho_{\mathbf{D}_k}}}(\mathbf{v}) = \left(\mathbf{A}^{\top} \mathbf{A} * \mathbf{B}_k^{\top} \mathbf{B}_k + \frac{\rho_{\mathbf{D}_k}}{2} \mathbf{I} \right)^{-1} \left(\text{Diag}(\mathbf{A}^{\top} \mathbf{X}_k \mathbf{B}_k) + \frac{\rho_{\mathbf{D}_k}}{2} \mathbf{v} \right), \quad (26)$$

where $\text{Diag}(\mathbf{A}^{\top} \mathbf{X}_k \mathbf{B}_k)$ is the vector containing the diagonal entries of $\mathbf{A}^{\top} \mathbf{X}_k \mathbf{B}_k$.

Algorithm 6: CMF-based ADMM updates for the \mathbf{A} -matrix

Result: $\mathbf{A}, \mathbf{Z}_{\mathbf{A}}, \mu_{\mathbf{A}}$

```

1 while stopping conditions are not met and max no. iterations not exceeded
  do
2    $\mathbf{A} \leftarrow \frac{\text{Equation (24)}}{\rho_{\mathbf{A}}} \text{prox}_{\frac{f_{\mathbf{A}}}{\rho_{\mathbf{A}}}}(\mathbf{Z}_{\mathbf{A}} - \mu_{\mathbf{A}})$ 
3    $\mathbf{Z}_{\mathbf{A}} \leftarrow \text{prox}_{\frac{g_{\mathbf{A}}}{\rho_{\mathbf{A}}}}(\mathbf{A} + \mu_{\mathbf{A}})$ 
4    $\mu_{\mathbf{A}} \leftarrow \mu_{\mathbf{A}} + \mathbf{A} - \mathbf{Z}_{\mathbf{A}}$ 
5 end

```

4.5 Stopping conditions and feasibility penalties

For the inner ADMM loops (Algorithms 4, 6, and 7), we follow [48] and use stopping conditions based on the primal and dual residuals:

$$\frac{\left\| \mathbf{x}^{(t,q)} - \mathbf{z}^{(t,q)} \right\|}{\left\| \mathbf{x}^{(t,q)} \right\|} \leq \epsilon, \quad \text{and} \quad \frac{\left\| \mathbf{z}^{(t,q)} - \mathbf{z}^{(t,q-1)} \right\|}{\left\| \mathbf{z}^{(t,q)} \right\|} \leq \epsilon. \quad (27)$$

In our case, \mathbf{z} is replaced with an auxiliary factor matrix (e.g., $\mathbf{Z}_{\mathbf{B}_k}$ or $\mathbf{Y}_{\mathbf{B}_k}$) and \mathbf{x} is replaced with the corresponding factor matrix (e.g., \mathbf{B}_k). The (t, q) superscript

Algorithm 7: CMF-based ADMM updates for the \mathbf{D}_k -matrices (C-matrix)

Result: $\mathbf{D}_k, \mathbf{Z}_{\mathbf{D}_k}, \boldsymbol{\mu}_{\mathbf{D}_k}$

- 1 **while** *stopping conditions are not met and max no. iterations not exceeded*
- do**
- 2 **for** $k \leftarrow 1$ **to** K **do**
- 3 $\mathbf{D}_k \xleftarrow{\text{Equation (26)}} \text{prox}_{\frac{f_{\mathbf{D}_k}}{\rho_{\mathbf{D}_k}}}(\mathbf{Z}_{\mathbf{D}_k} - \boldsymbol{\mu}_{\mathbf{D}_k})$
- 4 $\mathbf{Z}_{\mathbf{D}_k} \leftarrow \text{prox}_{\frac{g_{\mathbf{D}_k}}{\rho_{\mathbf{D}_k}}}(\mathbf{D}_k + \boldsymbol{\mu}_{\mathbf{D}_k})$
- 5 $\boldsymbol{\mu}_{\mathbf{D}_k} \leftarrow \boldsymbol{\mu}_{\mathbf{D}_k} + \mathbf{D}_k - \mathbf{Z}_{\mathbf{D}_k}$
- 6 **end**
- 7 **end**

represents the current outer and inner iteration number, respectively. For stopping the inner loops, we require that either all stopping conditions are fulfilled, or that a predefined number of iterations have been performed. In our experience, a low number (e.g., five) of inner iterations is sufficient.

For the outer loop, in Algorithm 2 we have two types of stopping conditions: loss decrease conditions and feasibility conditions. We stop once both the regularized sum of squared error and the relative feasibility gaps are below a given threshold, or their relative decrease is below a given threshold. That is, we ensure that

$$f^{(t)} + g^{(t)} < \epsilon^{\text{abs}} \quad \text{or} \quad \left| f^{(t-1)} + g^{(t-1)} - f^{(t)} - g^{(t)} \right| < \epsilon^{\text{rel}} \left(f^{(t-1)} + g^{(t-1)} \right),$$

where $f^{(t)}$ is the sum of squared errors after t outer iterations and $g^{(t)}$ is the sum of all regularization penalties after t outer iterations, are satisfied. We also ensure that

$$\mathbf{r}_{\mathbf{z}}^{(t)} = \left\| \mathbf{x}^{(t)} - \mathbf{z}^{(t)} \right\| < \epsilon^{\text{abs}} \quad \text{or} \quad |\mathbf{r}_{\mathbf{z}}^{(t)} - \mathbf{r}_{\mathbf{z}}^{(t-1)}| < \epsilon^{\text{rel}} \mathbf{r}_{\mathbf{z}}^{(t-1)},$$

is satisfied for all auxiliary variables \mathbf{z} . After exiting the AO-ADMM algorithm, it is important to verify that all primal feasibility gaps are sufficiently small. Otherwise, we may end up with components that violate the constraints.

In order to select the penalty parameters, we used the heuristic in [36], [37], setting

$$\begin{aligned} \rho_{\mathbf{A}} &= \frac{1}{R} \text{Tr} \left(\sum_{k=1}^K \mathbf{D}_k \mathbf{B}_k^{\top} \mathbf{B}_k \mathbf{D}_k \right) \\ \rho_{\mathbf{B}_k} &= \frac{1}{R} \text{Tr} \left(\mathbf{D}_k \mathbf{A}^{\top} \mathbf{A} \mathbf{D}_k \right), \\ \rho_{\mathbf{D}_k} &= \frac{1}{R} \text{Tr} \left(\mathbf{A}^{\top} \mathbf{A} * \mathbf{B}_k^{\top} \mathbf{B}_k \right), \end{aligned} \tag{28}$$

For convex regularization penalties, the ADMM-subproblems are guaranteed to converge with any ρ . However, for non-convex regularization penalties (e.g., the $\nu_{\mathcal{P}}$ in the \mathbf{B}_k updates) the value of ρ can affect ADMM's convergence properties [48].

4.6 Computational complexity of the AO-ADMM update steps

The computational complexities for the AO-ADMM update steps are given in Table 1³. To minimize the computational complexity, we compute parts of the right-hand side as well as the Cholesky factorization of the left hand side of all normal

³for the computational complexities of the CP-based updates, see supplementary Table SM1

Table 1: Computational complexities for the different AO-ADMM update steps. I , J and K denote the tensor size, R denotes the number of components and Q the number of iterations.

	Computational complexity
\mathbf{A}	$O(IJKR + JKR^2 + R^3 + IR^2Q)$
$\{\mathbf{D}_k\}_{k \leq K}$	$O(IJKR + IR^2 + JKR^2 + KR^3 + KR^2Q)$
$\{\mathbf{B}_k\}_{k \leq K}$	$O(IJKR + IR^2 + KR^3 + JKR^2Q)$

Table 2: Some proximable regularization penalties

Structure	Penalty	Proximal operator
Non-negativity	$\iota_{\mathbb{R}_+}$	$\text{prox}_{\iota_{\mathbb{R}_+}}(x) = \max(0, x)$
Graph Laplacian regularization	$\mathbf{x}^\top \mathbf{L} \mathbf{x}$	$\text{prox}_{\mathbf{x}^\top \mathbf{L} \mathbf{x}}(\mathbf{x}) = (\mathbf{L} + 0.5\mathbf{I})^{-1} \mathbf{x}$
TV regularization	$\sum_i x_i - x_{i-1} $	[52], [54]
Unimodality	$\iota_{\mathcal{U}}$	[53]
Unimodality and non-negativity	$\iota_{\mathcal{U} \cap \mathbb{R}_+}$	[46]
PARAFAC2 constraint	$\iota_{\mathcal{P}}$	Algorithm 5

equations (e.g. $\mathbf{X}_k^\top \mathbf{A} \mathbf{D}_k$ and the Cholesky factorization of $(\mathbf{D}_k \mathbf{A}^\top \mathbf{A} \mathbf{D}_k + \rho_{\mathbf{B}_k} \mathbf{I})$ for Equation (21)) only once per outer iteration and re-use it for the inner iterations.

4.7 Constraints

In this section, we review some background on proximal operators and list some useful proximable regularization penalties. If g is a proper lower semi-continuous convex function, the proximal operator, given by Equation (17), has a unique solution. For non-convex g , Equation (17) may not have a unique solution, and in those cases, we select one of the possible solutions. A large variety of regularization penalties admit closed form solutions or efficient algorithms for evaluating their proximal operator [50]–[53]. In this work, we evaluate the efficiency of our AO-ADMM scheme for four such penalty functions, given in Table 2. In particular, we impose non-negativity, graph Laplacian regularization, total variation (TV) regularization and unimodality. For more details on these penalties, see Section SM3 in the supplement.

4.7.1 The scale-indeterminacy of penalty-based regularization

Many penalty functions, such as graph Laplacian regularization and TV regularization, scale with the norm of the components. However, the PARAFAC2 loss function is invariant to scaling⁴. For example, \mathbf{A} can be multiplied with a constant, ϵ , so long as we multiply either all \mathbf{B}_k -matrices, or all \mathbf{D}_k -matrices by $1/\epsilon$. Thus, if we let $\epsilon \rightarrow 0$ we can obtain an arbitrarily small regularization penalty without affecting the components recovered by the algorithm (ignoring numerical difficulties arising as $\epsilon \rightarrow 0$).

To circumvent the scaling indeterminacy, we must regularize the norm of all factor matrices whenever the regularization for one factor matrix is norm-dependent. A straightforward way of regularizing the norm of the other factor matrices is with

⁴The observations in this section also hold for the CP decomposition.

ridge regularization, which we can incorporate into the proximal operator for the data-fidelity term. We will, for brevity’s sake, only show how the proximal operator for the \mathbf{A} -matrix is altered by this change (see supplementary for \mathbf{B}_k and \mathbf{D}_k alterations):

$$\text{prox}_{\frac{f_{\mathbf{A}} + \gamma \|\cdot\|_F^2}{\rho_{\mathbf{A}}}}(\mathbf{M}) = \left(\sum_{k=1}^K \mathbf{X}_k \mathbf{\Gamma}_k + \frac{\rho_{\mathbf{A}}}{2} \mathbf{M} \right) \left(\sum_{k=1}^K \mathbf{\Gamma}_k^T \mathbf{\Gamma}_k + \frac{2\gamma + \rho_{\mathbf{A}}}{2} \mathbf{I} \right)^{-1}, \quad (29)$$

where $\mathbf{\Gamma}_k = \mathbf{B}_k \mathbf{D}_k$ and γ is the ridge regularization penalty. Including ridge regularization this way does not necessarily increase the number of tunable parameters, which is apparent from the following theorem (proof in Section SM1):

Theorem 4.1. *Let $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ be a function satisfying $f(\mathbf{a}\mathbf{u}, a^{-2}\mathbf{v}) = f(\mathbf{u}, \mathbf{v})$, and let $r_u : \mathbb{R}^n \rightarrow \mathbb{R}$ and $r_v : \mathbb{R}^m \rightarrow \mathbb{R}$ be two (absolutely) homogeneous functions of degree d_u and d_v respectively. That is $r_u(\mathbf{a}\mathbf{u}) = |a|^{d_u} r_u(\mathbf{u})$ and $r_v(\mathbf{a}\mathbf{v}) = |a|^{d_v} r_v(\mathbf{v})$. Then the optimization problems*

$$\min_{\mathbf{u}, \mathbf{v}} [f(\mathbf{u}, \mathbf{v}) + a r_u(\mathbf{u}) + r_v(\mathbf{v})] \quad (30)$$

and

$$\min_{\mathbf{u}, \mathbf{v}} \left[f(\mathbf{u}, \mathbf{v}) + r_u(\mathbf{u}) + a^{2\frac{d_u}{d_v}} r_v(\mathbf{v}) \right] \quad (31)$$

are equivalent⁵ for any positive a .

Consequently, if we impose ridge regularization on two modes (e.g., \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$) with the same regularization parameter, and a homogeneous regularization penalty on the remaining mode (\mathbf{B}_k), then scaling the ridge penalty is equivalent to scaling the \mathbf{B}_k -penalty (set $\mathbf{u} = \{\mathbf{A}, \{\mathbf{D}_k\}_{k \leq K}\}$ and $\mathbf{v} = \{\mathbf{B}_k\}_{k \leq K}$ and apply Theorem 4.1).

5 Experiments

Here, we evaluate our proposed AO-ADMM scheme for constrained PARAFAC2 on a variety of experiments. We assess the performance in terms of accuracy (i.e., how well the underlying factors are captured), and computational efficiency on simulated datasets, and show the benefits of imposing constraints on the evolving mode of a PARAFAC2 model with real-world applications from chemometrics and neuroscience. Specifically, we use seven simulation setups. In Setup 1, 2, and 3, we demonstrate that the AO-ADMM scheme is both as accurate as and faster than the flexible coupling with HALS scheme for setups with non-negativity in all modes. Setup 1 also demonstrates the performance of the methods in the case of over- and under-estimation of the number of components. In Setup 3, we demonstrate that imposing unimodality constraints on $\{\mathbf{B}_k\}_{k \leq K}$ can improve accuracy and that the AO-ADMM scheme can be effective even when $\{\mathbf{B}_k\}_{k \leq K} \notin \mathcal{P}$. Next, in Setup 4 and 5, we respectively use a graph Laplacian- and TV-penalty to improve factor recovery. In Setup 6, we show that the cwSNR is an accurate predictor of factor recovery. Finally, in Setup SM1 in the supplementary, we compare the CP-based and CMF-based update steps on non-negative components, and show that they are similar in speed and accuracy. For Setup 1-6, we compare with a baseline based on PARAFAC2 ALS, which allows for constraints on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, but not on $\{\mathbf{B}_k\}_{k \leq K}$.

⁵Two optimization problems are equivalent if the solution to one can easily be transformed to the solution of the other [55].

5.1 Experimental setup

For both the PARAFAC2 AO-ADMM algorithm and the baselines, we used our Python implementations linked in the GitHub repository for the paper⁶ (There is also a MATLAB implementation linked in the same repository). The flexible coupling PARAFAC2 with HALS algorithm was implemented in Python closely following the MATLAB implementation by Cohen and Bro [30] (with some caching and reordering of computations for increased efficiency). To compute the proximal operator of the TV seminorm, we used the publicly available C implementation [54] of the improved direct TV denoising algorithm presented in [52]. Finally, the projection onto the set of non-negative unimodal vectors was implemented by thresholding the isotonic regression-vectors obtained with the algorithm from [53].

We set all stopping tolerances for the inner iterations to 10^{-5} with a maximum of 5 iterations (except for the unimodality setup, see Section 5.2.3). For the outer loop, we set all relative tolerances to 10^{-8} and all absolute tolerances to 10^{-7} .

For all experiments, we initialized the factor matrices, auxiliary matrices and scaled dual variables by drawing their elements from a uniform distribution between 0 and 1 (except the \mathbf{P}_k -matrices which were initialized as the first R columns of an identity matrix, thus ensuring that they are orthogonal). The same initial factor matrices were used for both the AO-ADMM experiments and the ALS experiments. We initialized the \mathbf{P}_k - and $\Delta_{\mathbf{B}}$ -matrices for ALS equally to the corresponding auxiliary variables for AO-ADMM. For all experiments, 20 initializations were used (except for Setup 6, where we used 50 initializations, and Setup 1 and SM1, where we used 10 initializations). We selected, for each model, the initialization that provided the lowest regularized sum of squared error among those that satisfied the stopping conditions (to ensure a low feasibility gap for the AO-ADMM components). If none satisfied the stopping conditions (i.e. all initializations reached the maximum number of iterations), we selected the initialization with the lowest regularized sum of squared error. All experiments imposed non-negativity on $\{\mathbf{D}_k\}_{k \leq K}$ to resolve the special sign indeterminacy of the PARAFAC2 model [44].

5.2 Simulations

For all simulation experiments, we first constructed simulated factor matrices. Then, after constructing a simulated tensor, \mathbf{X} , from those factor matrices (following Equation (4)), we added artificial noise as follows. First, we created a (possibly ragged) tensor \mathbf{E} with elements drawn from a normal distribution. Then, \mathbf{E} was scaled to have magnitude $\eta \|\mathbf{X}\|_F$ and added to \mathbf{X} following

$$\mathbf{X}_{\text{noisy}} = \mathbf{X} + \eta \|\mathbf{X}\|_F \frac{\mathbf{E}}{\|\mathbf{E}\|_F}. \quad (32)$$

Since this work introduces a scheme for regularizing the \mathbf{B}_k -matrices, we generated \mathbf{B}_k -matrices with structures tailored to various constraints. The elements of \mathbf{A} were drawn from a truncated normal distribution (i.e. $a_{ir} = \max(0, \check{a}_{ir})$, where $\check{a}_{ir} \sim \mathcal{N}(0, 1)$), and the elements of all \mathbf{D}_k -matrices were uniformly distributed ($\mathcal{U}(0.1, 1.1)$), avoiding near-zero elements, which can impede the recovery of \mathbf{B}_k .

To evaluate the different methods' performance on simulated data, we used the

⁶<https://github.com/MarieRoald/PARAFAC2-AOADM-SIMODS/>

relative sum of squared errors (Rel. SSE) given by

$$\text{Rel. SSE} = \frac{\sum_{k=1}^K \left\| \mathbf{X}_k - \mathbf{A} \mathbf{D}_k \mathbf{B}_k^\top \right\|_F^2}{\sum_{k=1}^K \left\| \mathbf{X}_k \right\|_F^2}, \quad (33)$$

and the factor match score (FMS), which is defined as

$$\text{FMS} = \frac{1}{R} \sum_{r=1}^R \left| \mathbf{a}_r^\top \hat{\mathbf{a}}_r \tilde{\mathbf{b}}_r^\top \hat{\tilde{\mathbf{b}}}_r \mathbf{c}_r^\top \hat{\mathbf{c}}_r \right|, \quad (34)$$

with all component vectors normalized and the superscript $\hat{\cdot}$ represents the estimated vectors. $\tilde{\mathbf{b}}_r$ and $\hat{\tilde{\mathbf{b}}}_r$ are vectors containing the concatenation of the r -th column of all \mathbf{B}_k -matrices and $\hat{\tilde{\mathbf{B}}}_k$ -matrices, respectively. Likewise, \mathbf{c}_r and $\hat{\mathbf{c}}_r$, contain the r -th diagonal entry of all \mathbf{D}_k -matrices and $\hat{\mathbf{D}}_k$ -matrices. The FMS ranges from zero to one where one indicates fully recovered components. Since the PARAFAC2 decomposition is only unique up to permutation and scaling, we find the optimal permutation of the components before computing the FMS. For the more difficult setups (Setup 2-3), we also evaluate the recovery of each factor matrix independently by computing $\text{FMS}_{\mathbf{A}}$, $\text{FMS}_{\mathbf{B}}$ and $\text{FMS}_{\mathbf{C}}$ only considering the relevant terms of Equation (34), i.e.:

$$\text{FMS}_{\mathbf{A}} = \frac{1}{R} \sum_{r=1}^R \left| \mathbf{a}_r^\top \hat{\mathbf{a}}_r \right|. \quad (35)$$

5.2.1 Setup 1: Shifting non-negative components

Data generation To compare the AO-ADMM scheme with the flexible coupling with HALS scheme in terms of speed and accuracy, we used a simulation setup inspired by the simulations in [30]. Specifically, we generated data by drawing the elements of $\hat{\mathbf{B}}$ from a truncated normal distribution. Then, to obtain \mathbf{B}_k , we shifted elements of $\hat{\mathbf{B}}$ by k indices, cyclically, setting $[\mathbf{B}_k]_{j,r} = \hat{\mathbf{B}}_{(j+k) \bmod J, r}$. This shifting yields non-negative \mathbf{B}_k -matrices that vary in a way that follows the PARAFAC2 constraint. Using this approach, we generated 50 different 3-component datasets, and 50 different 5-component datasets, each of size $30 \times 40 \times 50$, before adding noise with $\eta = 0.33$.

Experiment settings We used AO-ADMM and flexible coupling with HALS⁷ to decompose the data tensors with non-negativity imposed on all modes. We also compared with ALS with non-negativity on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. For the flexible coupling with HALS scheme, we used the same initialization of \mathbf{A} , $\{\mathbf{B}_k\}_{k \leq K}$, $\{\mathbf{D}_k\}_{k \leq K}$, $\{\mathbf{P}_k\}_{k \leq K}$ and $\mathbf{\Delta}_{\mathbf{B}}$ as the AO-ADMM scheme. All models were run until the stopping conditions were satisfied or a maximum of 2000 iterations was reached. For the 3-component data, we also fitted models with two and four components to compare the methods' robustness to over- and under-estimation of the number of components. For these experiments, we set the maximum number of iterations to 3000 and the absolute feasibility gap tolerance to 10^{-5} .

Results Figure 2 shows that the AO-ADMM scheme attains as good FMS as the flexible coupling with HALS scheme but faster. The AO-ADMM scheme was slightly

⁷We also ran experiments where we replaced HALS with other non-negative least squares algorithms [56], [57] and found that the relative performance of the AO-ADMM compared to the flexible coupling schemes did not change (results not shown).

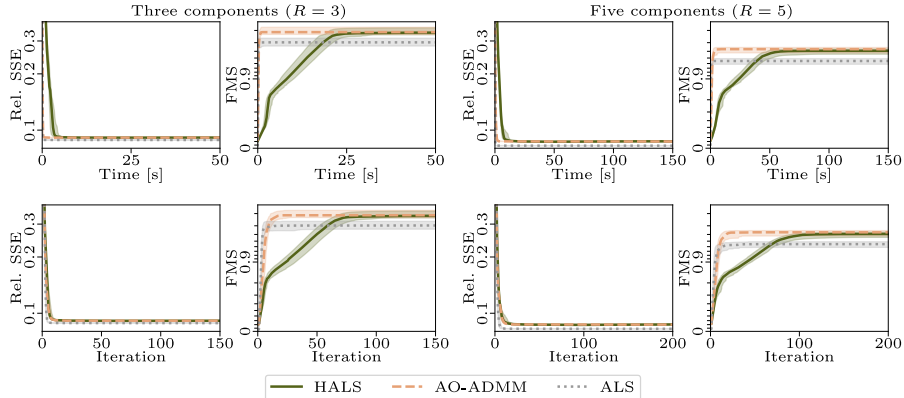


Figure 2: Setup 1: Diagnostic plots for the simulated data tensors. Median values are shown as solid lines, and the shaded area shows the interquartile range. The top row shows performance as a function of time, and the bottom row shows the performance as a function of iterations. The x-axes are cropped at 60 s ($R=3$) and 150 s ($R=5$) for the top row, and 150 iterations ($R=3$) and 200 iterations ($R=5$) for the bottom row. Subsequent iterations only decreased the feasibility gap (for AO-ADMM) and regularization penalty (for flexible coupling with HALS), not the FMS or relative SSE.

slower than ALS in terms of speed but got a higher FMS. In Figure 3, we see that the relative performance of the methods are the same in overfactoring and underfactoring with ALS getting lower FMS, and HALS being slower than AO-ADMM.

5.2.2 Setup 2: Non-negative components

Data generation Setup 1 focuses on shifting \mathbf{B}_k -matrices as in [30]. However, this implicitly assumes $\mathbf{B}_k = \mathbf{P}_k \mathbf{\Delta}_B$, with non-negative $\mathbf{\Delta}_B \in \mathbb{R}^{J \times R}$ and non-negative $\mathbf{P}_k \in \mathcal{O}_{J,J}$. Thus, the \mathbf{P}_k -matrices are assumed to have disjoint support. Here, we evaluate the different PARAFAC2 fitting algorithms on non-negative datasets without these implicit assumptions. Specifically, we construct a non-negative cross-product matrix, $\mathbf{X}^\top \mathbf{X}$, where $\mathbf{X} \in \mathbb{R}^{100 \times R}$ has elements from a truncated normal distribution and obtain the \mathbf{B}_k factor matrices by solving:

$$\begin{aligned} \min_{\mathbf{B}_k} & \left[\|\mathbf{B}_k^\top \mathbf{B}_k - \mathbf{X}^\top \mathbf{X}\|^2 \right] \\ \text{subject to} & \quad [\mathbf{B}_k]_{jr} \geq 0, \end{aligned} \quad (36)$$

using projected gradient descent with various random initializations for each \mathbf{B}_k (more details in the supplementary Subsection SM6.1). Following this, we constructed 50 three-component ragged tensors, consisting of 15 frontal slices, of size $50 \times J_k$, where J_k was a random integer between 50 and 100. We also evaluated non-negativity for all modes with increased collinearity in the \mathbf{C} -matrix, which we obtained by using $\tilde{\mathbf{D}}_k$ instead of \mathbf{D}_k , with $[\tilde{\mathbf{D}}_k]_{rr} = 0.5[\mathbf{D}_k]_{11} + 0.5[\mathbf{D}_k]_{rr}$. Finally, we added noise following (32) with 10 η -values logarithmically spaced between 0.5 and 2.5.

Experiment settings We decomposed each data tensor using PARAFAC2 AO-ADMM and HALS with non-negativity on all modes and PARAFAC2 ALS (with non-negativity on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$), both with a maximum of 6000 iterations.

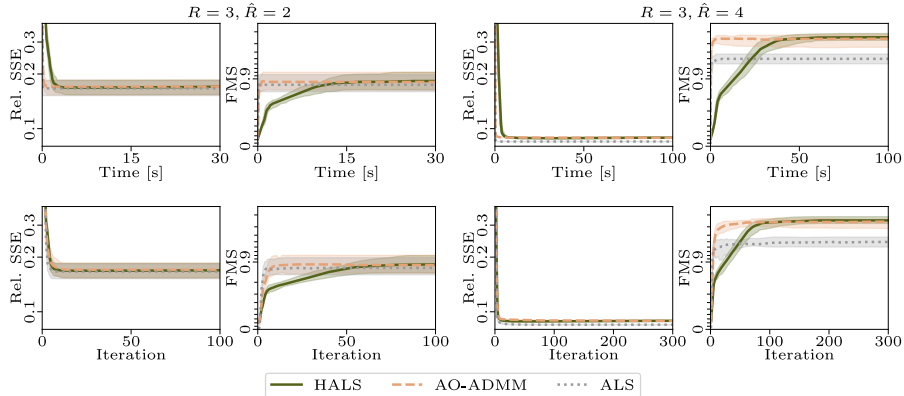


Figure 3: Setup 1: Diagnostic plots for the simulated data tensors with under- and over-estimation of the number of components. R represents the true number of components and \hat{R} represents the estimated number of components. To compute the FMS between two decompositions with different numbers of components, we use the subset of $\min(R, \hat{R})$ components that provides the highest FMS.

Results For some initializations, the ALS algorithm obtained degenerate solutions (solutions where two components are highly correlated in all modes but point in the opposite direction, and therefore cancel each other out) [58]. These initializations were excluded from the analysis. Furthermore, we also discarded the datasets where ALS gave degenerate solutions for all 20 initializations (these datasets were also removed for AO-ADMM)⁸. To identify degenerate solutions, we used the minimum triple-cosine (TC) based on the normalized component vectors [58]:

$$\text{TC} = \min_{r,s} \mathbf{a}_r^T \mathbf{a}_s [\mathbf{b}_1]_r^T [\mathbf{b}_2]_s \mathbf{c}_r^T \mathbf{c}_s, \quad (37)$$

excluding solutions with $\text{TC} < -0.85$ [5]. Figure 4 shows the resulting distributions of FMS for all setups with different noise levels. We observe that the fully constrained models obtained a higher FMS than the ALS algorithm for all setups. The largest improvement is for the $\text{FMS}_{\mathbf{B}}$ score, which is reasonable since the $\{\mathbf{B}_k\}_{k < K}$ components are the most challenging components to recover. However, the AO-ADMM and HALS models also improve the $\text{FMS}_{\mathbf{A}}$ compared to ALS (AO-ADMM improves the $\text{FMS}_{\mathbf{C}}$ as well). The improvement grows for higher noise levels, demonstrating that imposing non-negativity for all modes instead of just the non-evolving modes makes the model more robust to noise. Furthermore, as for Setup 1, HALS was slower in terms of both time and iterations (see Figures SM5 and SM6, respectively).

5.2.3 Setup 3: Unimodality constraints

Data generation For the unimodality constraints, we generated unimodal \mathbf{B}_k factor matrices similar to the approach in [31]. However, we also wanted to examine the performance for data that does not exactly follow the PARAFAC2 constraint. For this, we generated the r -th column of \mathbf{B}_k as the normal probability density function

$$[\mathbf{b}_k]_r = \mathbf{p}_{\text{normal}}(\mu_r + 0.41k, \sigma_{kr}), \quad (38)$$

⁸See supplementary Table SM4 for an overview of the number of datasets excluded because ALS gave degenerate solutions.

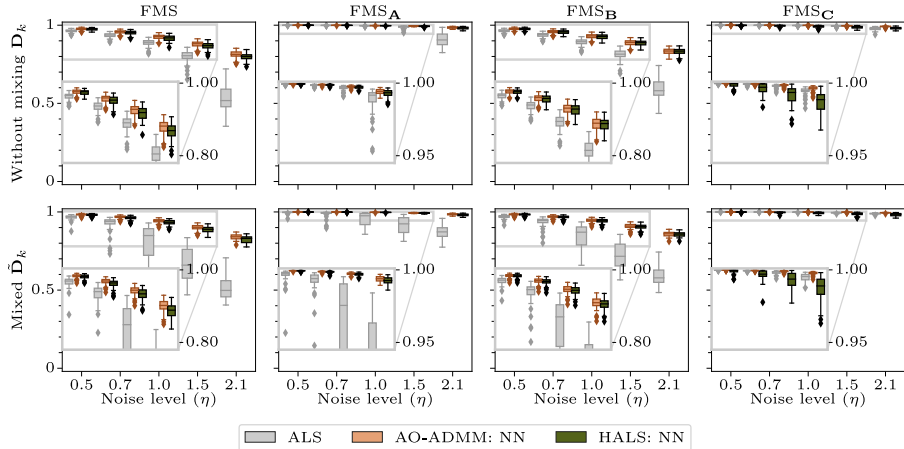


Figure 4: Setup 2: Boxplots showing the FMS for the different datasets. ALS refers to models fitted with ALS and non-negativity only on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, AO-ADMM: NN refers to models fitted with AO-ADMM with non-negativity imposed on all modes and HALS: NN refers to models fitted with HALS with non-negativity imposed on all modes. For the sake of space, this figure only shows every other noise level, see supplementary Figure SM4 for the remaining noise levels.

where $\mathbf{p}_{\text{normal}}(\mu, \sigma)$ is the PDF of a normal distribution sampled with 50 linearly spaced points between -10 and 10 and with given mean and standard deviation, $\mu_r \sim \mathcal{U}(-7, 0)$, $\sigma_{kr} = \sigma_r + n_{kr}$ with $\sigma_r \sim \mathcal{U}(0.5, 1)$ and $n_{kr} \sim \mathcal{N}(0, 0.1)$. The factor 0.41 was chosen since that is equivalent to a shift of one index. Using a Gaussian PDF gives non-negative, unimodal components and increasing μ for each slice produces shifting components, similar to setup 1. Furthermore, the random variation in σ_{kr} for each k leads to the components changing shape somewhat across slices in a way that will violate the PARAFAC2 constraint. Examining components that violate the PARAFAC2 constraint is interesting because real data might also not follow the PARAFAC2 constraint perfectly. The components are also highly collinear since the Gaussian curves are likely to overlap. Following this procedure, we constructed 50 datasets with five components, tensor sizes $10 \times 50 \times 15$ and $\eta = 0.33$. Subsection SM6.3 in the supplementary materials contains results from running the same experiment with $\sigma_{kr} = \sigma_r$ so the components follow the PARAFAC2 constraint.

Experiment settings For the constrained model, we imposed non-negativity on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, and unimodality as well as non-negativity on $\{\mathbf{B}_k\}_{k \leq K}$. As baselines, we used AO-ADMM and HALS with non-negativity imposed on all modes, and ALS with non-negativity imposed only on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. Each fitting algorithm ran until the stopping conditions were met or for a maximum of 2000 iterations.

During initial experiments, we observed that the feasibility gap ($\|\mathbf{B}_k - \mathbf{P}_k \Delta_{\mathbf{B}}\|_F$ and $\|\mathbf{B}_k - \mathbf{Z}_{\mathbf{B}_k}\|_F$) was too large for the AO-ADMM scheme to converge. This problem was mainly observed for the unimodality-constrained models, but it also occurred for the non-negativity constrained models. Therefore, we introduced three changes to the algorithm: We initialized the \mathbf{A} , $\{\mathbf{D}_k\}_{k \leq K}$, $\Delta_{\mathbf{B}}$ and \mathbf{P}_k -matrices by running the ALS algorithm for one iteration before AO-ADMM, increased the automatically selected penalty parameters for the evolving mode, $\rho_{\mathbf{B}_k}$ -s, by a factor of 10 and ran




Figure 5: Setup 3: Boxplots showing the FMS for different models fitted to datasets with unimodal \mathbf{B}_k -matrices. ALS represents models fitted with non-negativity imposed on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, HALS and NN represents models fitted with non-negativity imposed on all modes using HALS and AO-ADMM respectively, and NN&U represents models fitted with non-negativity on all modes and unimodality imposed on $\{\mathbf{B}_k\}_{k \leq K}$ using AO-ADMM.

the inner (ADMM) iterations for at most 20 iterations instead of 5.

Results As with Setup 2, we observed that the ALS scheme sometimes provided degenerate solutions, which were disregarded using the same heuristic. Datasets where all initializations fitted with ALS were degenerate were excluded from analysis, leaving 43 out of 50 datasets. Figure 5 shows that models fitted with AO-ADMM obtained a higher FMS than those fitted with ALS and HALS⁹. Furthermore, imposing non-negativity and unimodality on $\{\mathbf{B}_k\}_{k \leq K}$ improved the FMS compared to only imposing non-negativity and substantially improved FMS compared to imposing no constraints on $\{\mathbf{B}_k\}_{k \leq K}$. Also, we observed an increase in the recovery of all factor matrices for models fitted with unimodality constraints compared to those fitted without. In Figure 6, we see the true and estimated factors for one of the datasets and the cross product of \mathbf{B}_k for each slice, k , in this dataset is shown in Figure 7. We see from the figure that the cross product is not constant. However, despite this, the components are recovered, which demonstrates that the PARAFAC2 AO-ADMM scheme works well even with slight violations of the PARAFAC2 constraint.

5.2.4 Setup 4: Graph Laplacian regularization

Data generation To assess graph Laplacian regularization, we simulated *smooth* components, using a simulation setup inspired by [29]. For this setup, we first select a D -dimensional space of smooth vectors, \mathcal{V} , and construct an orthonormal matrix \mathbf{M} whose columns span \mathcal{V} . Next, we observe that if we have a collection of matrices, $\{\tilde{\mathbf{B}}_k\}_{k \leq K} \in \mathcal{P}$, then we can multiply each factor matrix, $\tilde{\mathbf{B}}_k$, with \mathbf{M} , thus obtaining a new collection of matrices, $\{\mathbf{B}_k\}_{k \leq K} = \{\mathbf{M}\tilde{\mathbf{B}}_k\}_{k \leq K} \in \mathcal{P}$, whose columns are smooth.

We constructed \mathbf{M} from a basis for cubic polynomials on the interval $[-1, 1]$, orthogonalized using the SVD. To generate random $\{\tilde{\mathbf{B}}_k\}_{k \leq K} \in \mathcal{P}$, we set $\tilde{\mathbf{B}}_k = \mathbf{P}_k \Delta_{\mathbf{B}}$, drawing the elements of $\Delta_{\mathbf{B}}$ from a standard normal distribution and generated the \mathbf{P}_k -matrices as random orthogonal matrices (obtained from the QR factorization: $\mathbf{P}_k \mathbf{R} = \mathbf{S}_k$, where $[\mathbf{S}_k]_{ij} \sim \mathcal{N}(0, 1)$). Following this procedure, we set $R = 3$ and constructed 20 different $30 \times 200 \times 30$ simulated data tensors with $\eta = 0.5$.

Experiment settings Since the components are smooth, it is reasonable to penalize local differences. Thus, we constructed the similarity graph, such that neighbouring indices in the vectors had a similarity score of 1, and all non-neighbouring indices of the vectors had a similarity score of 0. To decide the strength of the regularization

⁹For timing plots, see Figure SM7 in the supplement.

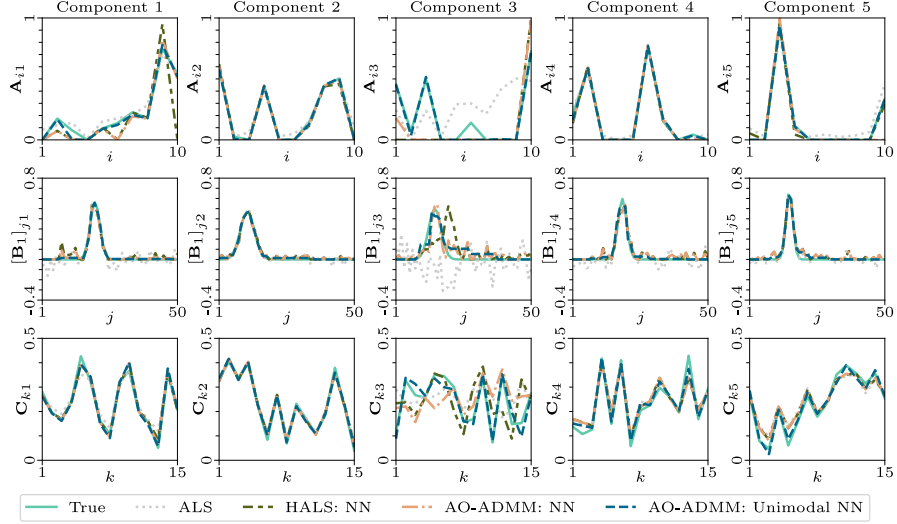


Figure 6: Setup 3: Plots showing the true and estimated components for one of the datasets. AO-ADMM: NN and HALS: NN represents the model fitted with non-negativity imposed on all modes using AO-ADMM and HALS, respectively, Unimodal NN represents the model fitted with non-negativity on all modes and unimodality imposed on $\{\mathbf{B}_k\}_{k \leq K}$ using AO-ADMM, and ALS represents the model fitted with non-negativity imposed on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. The cwSNRs (left to right) are $-4.4, -2.6, -12, -11$ and -4.1 dB for this frontal slice.




Figure 7: Setup 3: Plot showing how the cross product, $\Phi(k) = \mathbf{B}_k^T \mathbf{B}_k$, changes across frontal slices for the same dataset as Figure 6.

penalty, we conducted a grid-search of four regularization parameters logarithmically spaced between 1 and 1000. Also, following Theorem 4.1, we used ridge regularization on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, setting the regularization strength to 0.1. We compared with models without ridge regularization to show that a penalty-based regularization on one mode requires penalizing the norm of the remaining factor matrices. In addition, we compared with an ALS baseline without regularization. All models were ran until the stopping conditions were met or they reached 5000 iterations.

Results From Figures 8 and 9, we see that imposing smoothness regularization yielded smoother components and a higher FMS compared to using ALS without regularization. Furthermore, ridge regularization on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ was essential for obtaining improvement with graph Laplacian regularization on the \mathbf{B}_k -matrices, which is expected from the observations in Section 4.7.1. To ensure that this was not an artifact from the initialization selection scheme, we also compared the initializations that obtained the highest FMS for each dataset, which showed the same behavior (see supplement, Figure SM12). Another notable observation is that none of the models fitted without ridge regularization on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, but with regularization on \mathbf{B}_k converged (see supplement, Table SM5). Moreover, when we imposed regularization on only the \mathbf{B}_k -matrices, the regularization penalty (when calculated after normalizing the \mathbf{B}_k -matrices) decreased only in early iterations before approaching the same value as the unregularized (ALS) model. Figure SM10 in the supplement shows a more detailed visualization. This behavior demonstrates that if we use penalty based regularization on the \mathbf{B}_k -matrices without penalizing the norm of the \mathbf{A} -matrix and \mathbf{D}_k -matrices, then the regularization will have little effect.

5.2.5 Setup 5: Total variation regularization

Data generation For assessing TV regularization, we simulated piecewise constant components for ragged tensors. To construct these components, we used the following scheme: For each \mathbf{B}_k -matrix, we generated a random binary matrix with orthogonal columns, $\mathbf{Q}_k \in \mathbb{R}^{J_k \times 4}$, and at most two discontinuities per column (i.e., piecewise constant). Then, we normalized these matrices, obtaining orthonormal matrices, \mathbf{Q}_k , and set $\mathbf{B}_k = \mathbf{Q}_k \mathbf{\Omega}$, where $\mathbf{\Omega} \in \mathbb{R}^{4 \times R}$ has two non-zero elements per column, drawn from a standard normal distribution. This scheme led to piecewise constant components with at most four jumps. We set $R = 3$ and constructed 20 such simulated ragged data tensors with 30 frontal slices of size $30 \times J_k$, with J_k drawn uniformly between 200 and 250. Following Equation (32), we added noise with $\eta = 0.5$.

Experiment settings We imposed TV regularization on the \mathbf{B}_k -matrices and ridge regularization on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. Setting the ridge regularization parameter to both 0 and 0.1, we performed a grid search for the TV regularization parameter with five logarithmically spaced values between 0.001 and 10. As baseline we used ALS without regularization on $\{\mathbf{B}_k\}_{k \leq K}$. All methods were run until the stopping conditions were fulfilled or they reached 5000 iterations.

Results Figure 8 shows that the TV-regularized PARAFAC2 AO-ADMM algorithm performs better than the unregularized ALS algorithm for most regularization strengths. Moreover, from Figure 9 we see that the TV regularized factors are piecewise constant when ridge was imposed on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. Similar to setup 4, we observe that for AO-ADMM, none of the initializations converged without ridge regularization on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ (see supplement, Table SM6) and the regularization

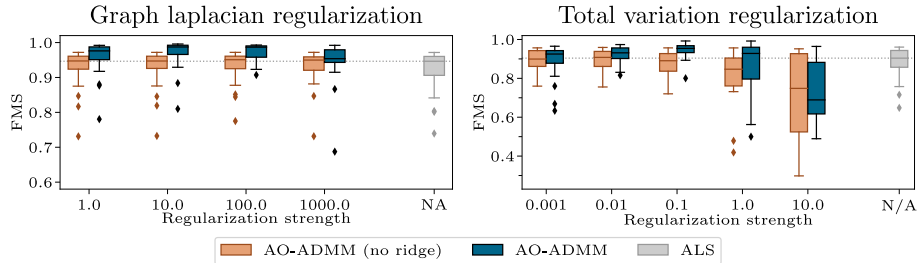


Figure 8: Setup 4 and 5: Boxplots showing the FMS obtained with different levels of regularization.

(when calculated after normalizing the \mathbf{B}_k -matrices) decreased only in early iterations before increasing (see supplement Figure SM11).

5.2.6 Setup 6: Evaluating the cwSNR

Dataset generation To evaluate the cwSNR ability to assess the recovery of the \mathbf{B}_k components, we simulated five different 5-component datasets of size $30 \times 40 \times 100$. For the \mathbf{B}_k factor matrices, we drew the elements of \mathbf{B}_1 from a truncated normal distribution and then shifted them cyclically by k entries. After constructing the data tensors, we added noise following Equation (32) with $\eta = 0.1$, $\eta = 0.33$ and $\eta = 0.5$.

Experiment setup Each model was ran for 2000 iterations or until the stopping conditions were satisfied. Then, to assess the degree of recovery for a given $[\hat{\mathbf{b}}_k]_r$, we used the cosine similarity (SIM), given by (for normalized component vectors)

$$\text{SIM}_{kr} = [\mathbf{b}_k]_r^\top [\hat{\mathbf{b}}_k]_r. \quad (39)$$

Results Figure 10 shows the results from these experiments. The top row shows the cosine similarity scores obtained with non-negativity on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ (ALS), the bottom row shows the recovery scores obtained with non-negativity on all modes (AO-ADMM) and the columns correspond to the different noise levels. We see that the expected recovery score grows monotonically with the cwSNR for both models.

5.3 Neuroimaging application

For neuroimaging applications, our motivation for imposing constraints on PARAFAC2 components is to improve their interpretability. Previous work has shown that PARAFAC2 can reveal dynamic networks (spatial dynamics) from fMRI data arranged as a *subjects by voxels by time windows* tensor [18]. However, the model is prone to noise affecting the extracted networks (see Figure 11). To investigate if smoothness inducing regularization can improve interpretability, we used images from the MCIC collection [59], which contains fMRI-scans from healthy controls and patients with schizophrenia. We used the sensory motor (SM) task data, with the same feature extraction and preprocessing as [18].

We analyzed the fMRI data tensor using both regularized PARAFAC2 fitted with AO-ADMM and unregularized PARAFAC2 fitted with ALS, both ran for at most 8000 iterations. The \mathbf{A} , \mathbf{B}_k and $\{\mathbf{D}_k\}_{k \leq K}$ factor matrices represented the subject-mode,

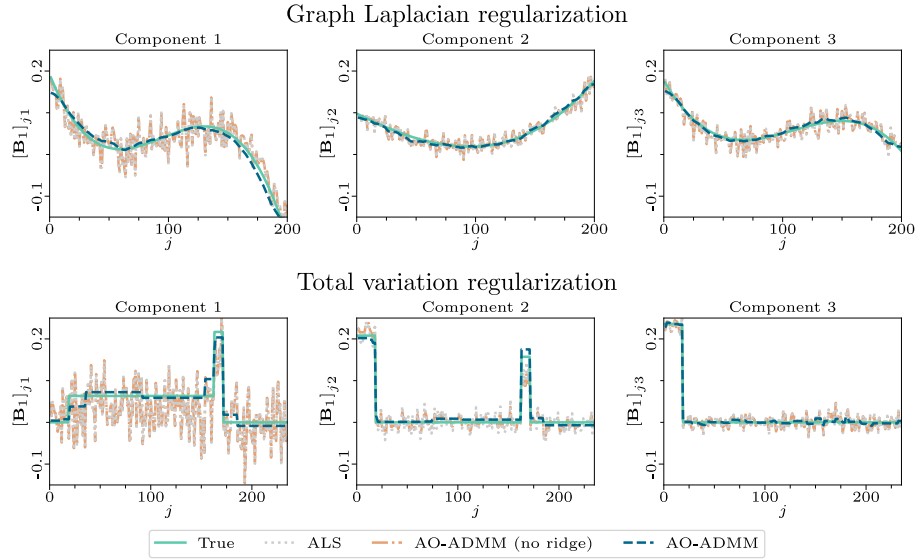


Figure 9: Setup 4 and 5: Plots showing true and estimated components. On the top, we see one of the datasets where graph Laplacian regularization was used to obtain smooth components and on the bottom we see another dataset where TV regularization was used to obtain piecewise constant components. The cwSNRs for the top row are $-8.6, 0.031$ and -3.7 , respectively. For the bottom row, the cwSNRs are $-14, -5.6$ and 0.36 , respectively.




Figure 10: Setup 6: Scatter plot showing factor recovery (SIM) plotted against cwSNR for models fitted with non-negativity on all modes (AO-ADMM) and non-negativity only on \mathbf{A} and \mathbf{C} (ALS). Each dot represents the cosine similarity score for one column of one \mathbf{B}_k factor matrix for one dataset (all five datasets are aggregated in each plot).

voxel-mode and time-mode components, respectively. This configuration allows the voxel-mode components that indicate activation networks to evolve over time. To reduce the noise in the components, we chose a graph Laplacian regularizer based on the image gradient, thus penalizing large differences in neighboring voxels. We imposed ridge regularization on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ to resolve the scaling indeterminacy, and non-negativity constraints on $\{\mathbf{D}_k\}_{k \leq K}$ to resolve the sign indeterminacy.

The spatial regularizer encourages similar activation for neighboring voxels, which is a reasonable assumption for fMRI images [60]. Specifically, we used graph Laplacian regularization where $w_{ij} = 1$ if voxel i and j are neighboring voxels (in the von Neumann sense) and $w_{ij} = 0$ otherwise. The proximal operator of this penalty function involves solving a large linear system, similar to that obtained when solving Helmholtz’ equation with a finite difference scheme. To solve this system, we, therefore, used the conjugate gradient method, with a smoothed aggregation preconditioner, setting the near-nullspace component to a vector consisting only of ones.

We fixed the ridge penalty to 0.1, and performed a grid-search for the optimal smoothness penalty. The selected regularization parameter, 10, provided the strongest smoothing-effect while also not deteriorating the interpretability of the model.

Figure 11 shows a comparison of the unregularized and the regularized model. We see that imposing regularization yields smoother, more interpretable, components. The regularized sensorimotor component (i.e., component 1) shows the same significantly stronger activation in healthy controls compared to patients as the unregularized component (the p -values obtained using a two-sample t-test on the corresponding subject-mode component vector is 2.29×10^{-4} for AO-ADMM and 2.98×10^{-4} for ALS) while the regions of activation are less affected by noise, and thus the interpretation is clearer.

5.4 Chemometrics application

Another application where PARAFAC2 is well-suited is the analysis of gas chromatography mass spectrometry (GC-MS) measurements [14], [61]. In GC-MS measurements, a sample is injected into a column that lets different molecules pass through it at different velocities depending amongst others on their affinity to the column (GC). At distinct time points at the end of the column, the mass spectrum is measured (MS). From this, we obtain a three-way tensor where one mode represents the mass spectrum, one mode represents the time after injecting the sample (retention time), and one mode represents the different samples. When we analyze the GC-MS data with PARAFAC2, we obtain components that represent the mass spectrum (\mathbf{A}) and elution profile ($\{\mathbf{B}_k\}_{k \leq K}$) for the different chemical compounds as well as their concentration in different samples ($\{\mathbf{D}_k\}_{k \leq K}$).

The dataset used here stems from a project on fermentation of apple wine using different microorganisms. The data is typical for untargeted chemical profiling as commonly done in flavor research. Samples were taken every eight hours from the headspace of the fermentation tank throughout the fermentation and measured on an Agilent 5973 MSD, Agilent Technologies, Santa Clara, USA.

Gathering GC-MS measurements can be time-consuming, and it is therefore of interest to enable the analysis of smaller GC-MS datasets. The quality of the analysis may depend on the number of samples, K . For this dataset, if we analyze all of the samples, constraints on \mathbf{B}_k are not needed for recovery and PARAFAC2 with ALS is, therefore, sufficient. However, constraining all modes may reduce the number of samples required for accurately capturing the underlying components. To evaluate whether constraining $\{\mathbf{B}_k\}_{k \leq K}$ reduces the number of samples needed, we constructed

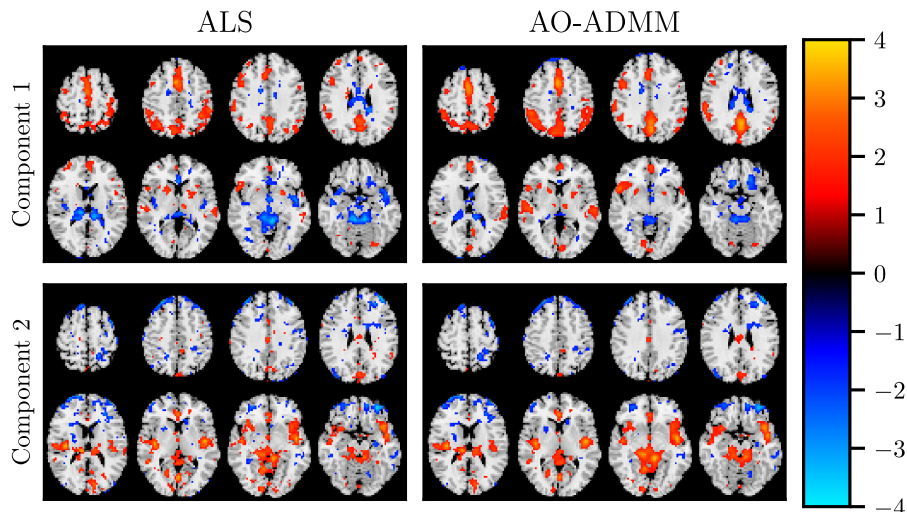


Figure 11: Voxel-mode components of the two-component unregularized (left) and regularized (right) PARAFAC2 models fitted to the fMRI data. The composite plots are generated by z-scoring the components and only showing voxels with activation > 1.5 . The regularization penalties (after normalization) for the voxel-mode components are 10.4 and 17.6 for the regularized and unregularized model, respectively, which demonstrates that the regularization quantifiably increases smoothness. We see that the regularized model has less speckle noise and more defined regions of activation. Furthermore, the regularized components are more focal (i.e., have high intensity in the mid-parts), which is expected of spatial fMRI maps, and thus they lend themselves to better association with physical quantities of interest. For the component that shows a statistically significant difference between the two groups (component 1), we see activation in the primary and secondary motor and cerebellum, and also auditory cortex, which are regions expected to be activated by the SM task. This activation is less noisy and easier to discern for the regularized model.

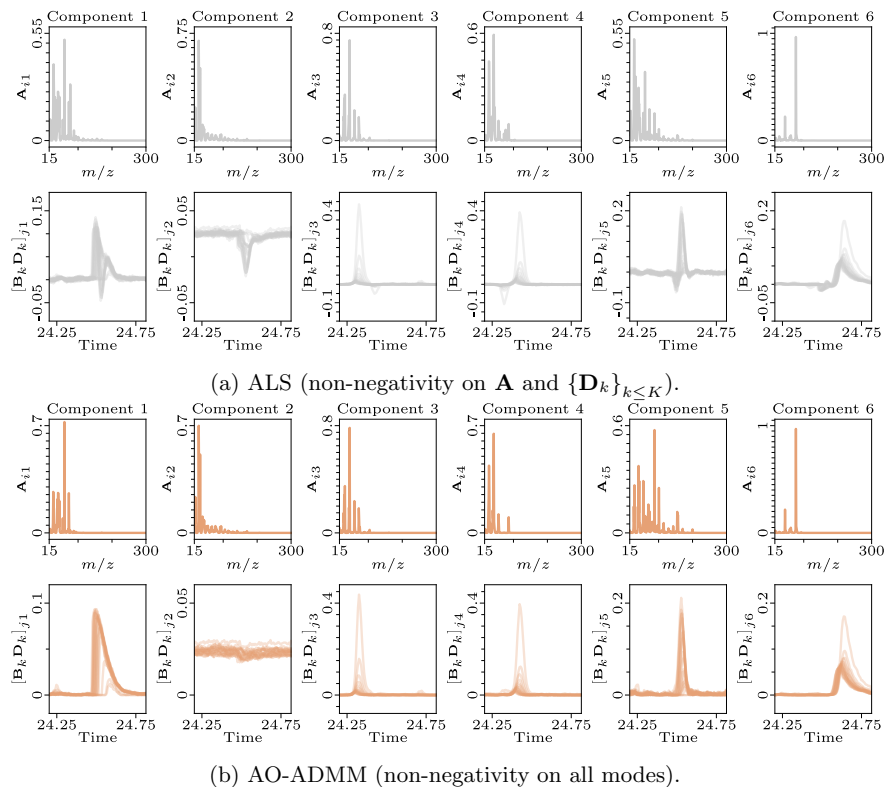


Figure 12: Mass spectra (top) and elution profiles scaled by concentration (bottom). For both models, the second component represents noise and the other components represent chemical information.

the data tensor from only a third of the 57 samples (i.e., the first 19 samples), forming a $286 \times 95 \times 19$ tensor. We then analyzed the GC-MS tensor using PARAFAC2 fitted with both AO-ADMM and ALS, both ran for at most 6000 iterations. As for the simulations, we ran 20 random initializations, selecting the runs with the lowest SSE.

PARAFAC2 is well suited because it allows the elution profiles to differ across different samples. Thus, PARAFAC2 can account for retention shifts and other shape changes that can occur between different samples containing the same chemical. Neither concentrations, elution profiles nor mass spectra should contain negative entries, so we impose non-negativity constraints on all factor matrices for PARAFAC2 with AO-ADMM and on the \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$ factor matrices for PARAFAC2 with ALS.

Figure 12 shows the 6-component models obtained with ALS and AO-ADMM. We see that both models include one noise component and five components that contain chemical information. However, the model fitted without non-negativity on \mathbf{B}_k yields unphysical elution profiles (negative peaks and multiple peaks), which demonstrates the need for more samples if we do not impose constraints on $\{\mathbf{B}_k\}_{k \leq K}$. For the model fitted with non-negativity on all modes, these phenomena are much less apparent, and we only observe minor secondary peaks, particularly for low-concentration samples. Thus, constraining the \mathbf{B}_k -matrices can improve the interpretative value of the PARAFAC2 components and reduce the number of samples required for analysis.

6 Discussion

In this paper, we have proposed an efficient algorithmic framework based on AO-ADMM that allows us to fit PARAFAC2 models with any proximable regularization penalty on all modes. Furthermore, we show that our approach can successfully apply hard constraints, e.g., non-negativity and unimodality, or penalty based regularization such as TV and graph Laplacian regularization on $\{\mathbf{B}_k\}_{k \leq K}$. Our numerical experiments show that the AO-ADMM framework is accurate, and faster than the flexible coupling approach for non-negative PARAFAC2 and that their relative performance was unchanged with under- and over-estimation of the rank.

With both simulated data and real applications, we demonstrated that constraining the \mathbf{B}_k -matrices, can improve recovery and interpretability. In particular, constraints can be necessary to recover the components accurately from noisy data when the component vectors are highly collinear. We also showed that the proposed scheme recovered the components even in cases where the PARAFAC2 constraint is slightly violated. Furthermore, we demonstrated that all modes should be regularized when the regularization penalty of one mode is dependent on the norm.

Nevertheless, there are possibilities for improvement. For example, the ADMM update steps are based on the projection $\text{prox}_{\mathcal{P}}$, for which we proposed an efficient approximation algorithm based on AO. However, the mathematical properties of \mathcal{P} are, to the best of our knowledge, still not well understood, and future analysis may improve upon this projection. Another possible extension is support for other loss functions, such as the KL-divergence or weighted least squares. Our scheme could also be extended to other models such as block term decomposition 2 (BTD2) [62], which is a four-way extension of the block term decomposition (BTD) [63] with evolving factor matrices that follow the PARAFAC2 constraint. AO-ADMM has already been used to fit BTD models [64], and our work could be straightforwardly adapted for fitting constrained BTD2 models with AO-ADMM. It may also be beneficial to use Nesterov-type extrapolation, which has been successfully used to speed up ALS schemes for CP [65], [66] as well as ALS and flexible coupling schemes for PARAFAC2 [67].

Acknowledgments

We would like to thank Jesper Løve Hinrich and Remi Cornillet for insightful discussions on constrained PARAFAC2. This work has benefited from the Experimental Infrastructure for Exploration of Exascale Computing (eX3), which is supported by the Research Council of Norway under contract 270053.

References

- [1] M. Roald, C. Schenker, J. E. Cohen, and E. Acar, “PARAFAC2 AO-ADMM: Constraints in all modes,” in *EUSIPCO’21: Proc. 2021 29th Eur. Signal Process. Conf.*, EURASIP, 2021.
- [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, ISSN: 1860-0980.
- [3] R. A. Harshman, “Foundations of the PARAFAC procedure: Models and conditions for an “explanatory” multi-modal factor analysis,” *UCLA working papers in phonetics*, vol. 16, pp. 1–84, 1970.

- [4] F. L. Hitchcock, “The expression of a tensor or a polyadic as a sum of products,” *J. Math. Phys.*, vol. 6, no. 1, pp. 164–189, 1927.
- [5] R. Bro, “PARAFAC. tutorial and applications,” *Chemom. Intel. Lab. Syst.*, vol. 38, no. 2, pp. 149–172, 1997.
- [6] M. Mørup, L. K. Hansen, C. S. Herrmann, J. Parnas, and S. M. Arnfred, “Parallel factor analysis as an exploratory tool for wavelet transformed event-related EEG,” *NeuroImage*, vol. 29, no. 3, pp. 938–947, 2006.
- [7] E. Acar, C. A. Bingol, H. Bingol, R. Bro, and B. Yener, “Multiway analysis of epilepsy tensors,” *Bioinformatics*, vol. 23, no. 13, pp. i10–i18, 2007.
- [8] A. H. Williams, T. H. Kim, F. Wang, S. Vyas, S. I. Ryu, K. V. Shenoy, M. Schnitzer, T. G. Kolda, and S. Ganguli, “Unsupervised discovery of demixed, low-dimensional neural dynamics across multiple timescales through tensor component analysis,” *Neuron*, vol. 98, no. 6, pp. 1099–1115, 2018.
- [9] D. M. Dunlavy, T. G. Kolda, and E. Acar, “Temporal link prediction using matrix and tensor factorizations,” *ACM TKDD*, vol. 5, no. 2, Article 10, 2011.
- [10] T. G. Kolda and B. W. Bader, “Tensor decompositions and applications,” *SIAM Rev.*, vol. 51, no. 3, pp. 455–500, 2009.
- [11] E. Acar and B. Yener, “Unsupervised multiway data analysis: A literature survey,” *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 1, pp. 6–20, Jan. 2009.
- [12] E. E. Papalexakis, C. Faloutsos, and N. D. Sidiropoulos, “Tensors for data mining and data fusion: Models, applications, and scalable algorithms,” *ACM Trans. Intell. Syst. Technol.*, vol. 8, no. 2, Article 16, 2016.
- [13] R. A. Harshman, “PARAFAC2: Mathematical and technical notes,” *UCLA working papers in phonetics*, vol. 22, pp. 30–44, 1972.
- [14] R. Bro, C. A. Andersson, and H. A. L. Kiers, “PARAFAC2 - Part II. Modeling chromatographic data with retention time shifts,” *J. Chemom.*, vol. 13, no. 3-4, pp. 295–309, 1999, ISSN: 0886-9383.
- [15] A. Afshar, I. Perros, E. E. Papalexakis, E. Searles, J. Ho, and J. Sun, “COPA: Constrained PARAFAC2 for Sparse & Large Datasets,” in *CIKM’18: Proc. ACM Int. Conf. Inf. Knowl. Manag.*, 2018, pp. 793–802.
- [16] P. A. Chew, B. W. Bader, T. G. Kolda, and A. Abdelali, “Cross-language information retrieval using PARAFAC2,” in *KDD’07: Proc. 13th ACM SIGKDD Int. Conf. Knowl. Discov. and Data Mining*, 2007, pp. 143–152.
- [17] K. Madsen, N. Churchill, and M. Mørup, “Quantifying functional connectivity in multi-subject fMRI data using component models,” *Human Brain Mapping*, vol. 38, no. 2, pp. 882–899, 2017.
- [18] M. Roald, S. Bhinge, C. Jia, V. Calhoun, T. Adali, and E. Acar, “Tracing network evolution using the PARAFAC2 model,” in *ICASSP’20: Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2020.
- [19] I. Lehmann, E. Acar, T. Hasija, M. A. Akhonda, V. D. Calhoun, P. J. Schreier, and T. Adali, “Multi-task fMRI data fusion using IVA and PARAFAC2,” in *ICASSP’22: Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2022.
- [20] D. Nion and N. D. Sidiropoulos, “Adaptive algorithms to track the PARAFAC decomposition of a third-order tensor,” *IEEE Trans. Signal Process.*, vol. 57, pp. 2299–2310, 2009.

- [21] M. Vandecappelle, N. Vervliet, and L. D. Lathauwer, “Nonlinear least squares updating of the canonical polyadic decomposition,” in *EUSIPCO’17: Proc. 25th Eur. Signal Proc. Conf.*, EURASIP, 2017, pp. 663–667.
- [22] M. Mardani, G. Mateos, and G. B. Giannakis, “Subspace learning and imputation for streaming big data matrices and tensors,” *IEEE Trans. Signal Process.*, vol. 63, no. 10, pp. 2663–2677, 2015.
- [23] T. Maehara, K. Hayashi, and K. Kawarabayashi, “Expected tensor decomposition with stochastic gradient descent,” in *AAAI’16: Proc. 30th AAAI Conf. Artif. Intell.*, 2016, pp. 1919–1925.
- [24] T. G. Kolda and D. Hong, “Stochastic gradients for large-scale tensor decomposition,” *SIAM J. Math. Data Sci.*, vol. 2, no. 4, pp. 1066–1095, 2020.
- [25] H. A. L. Kiers, J. M. F. Ten Berge, and R. Bro, “PARAFAC2 - Part I. A direct fitting algorithm for the PARAFAC2 model,” *J. Chemom.*, vol. 13, no. 3-4, pp. 275–294, 1999.
- [26] Y.-X. Wang and Y.-J. Zhang, “Nonnegative matrix factorization: A comprehensive review,” *IEEE Trans. Knowl. Data Eng.*, vol. 25, no. 6, pp. 1336–1353, 2012.
- [27] R. Bro and S. De Jong, “A fast non-negativity-constrained least squares algorithm,” *J. Chemom.*, vol. 11, no. 5, pp. 393–401, 1997.
- [28] M. P. Friedlander and K. Hatz, “Computing non-negative tensor factorizations,” *Optim. Methods Softw.*, vol. 23, no. 4, pp. 631–647, 2008.
- [29] N. E. Helwig, “Estimating latent trends in multivariate longitudinal data via Parafac2 with functional and structural constraints,” *Biom. J.*, vol. 59, no. 4, pp. 783–803, 2017.
- [30] J. E. Cohen and R. Bro, “Nonnegative PARAFAC2: A flexible coupling approach,” in *LVA/ICA’18*, 2018, pp. 89–98.
- [31] M. H. Van Benthem, T. J. Keller, G. D. Gillispie, and S. A. DeJong, “Getting to the core of PARAFAC2, a nonnegative approach,” *Chemom. Intel. Lab. Syst.*, vol. 206, p. 104 127, 2020.
- [32] K. Yin, A. Afshar, J. C. Ho, W. K. Cheung, C. Zhang, and J. Sun, “LogPar: Logistic PARAFAC2 factorization for temporal binary data with missing values,” in *KDD’20: Proc. 26th ACM SIGKDD Int. Conf. Knowl. Discov. and Data Mining*, 2020, pp. 1625–1635.
- [33] J. D. Carroll, S. Pruzansky, and J. B. Kruskal, “Candelinec: A general approach to multidimensional analysis of many-way arrays with linear constraints on parameters,” *Psychometrika*, vol. 45, no. 1, pp. 3–24, 1980.
- [34] R. Bro and C. A. Andersson, “Improving the speed of multiway algorithms: Part II: Compression,” *Chemom. Intel. Lab. Syst.*, vol. 42, no. 1-2, pp. 105–113, 1998.
- [35] H. A. L. Kiers, “A three-step algorithm for CANDECOMP/PARAFAC analysis of large data sets with multicollinearity,” *J. Chemom.*, vol. 12, no. 3, pp. 155–171, 1998.
- [36] K. Huang, N. D. Sidiropoulos, and A. P. Liavas, “A flexible and efficient algorithmic framework for constrained matrix and tensor factorization,” *IEEE Trans. Signal Process.*, vol. 64, no. 19, pp. 5052–5065, 2016.

- [37] C. Schenker, J. E. Cohen, and E. Acar, “A flexible optimization framework for regularized matrix-tensor factorizations with linear couplings,” *IEEE J. Sel. Top. Signal Process.*, vol. 15, no. 3, pp. 506–521, 2021.
- [38] C. Schenker, J. E. Cohen, and E. Acar, “An optimization framework for regularized linearly coupled matrix-tensor factorization,” in *EUSIPCO’20: Proc. 2020 28th Eur. Signal Process. Conf.*, EURASIP, 2021, pp. 985–989.
- [39] Y. Ren, J. Lou, L. Xiong, and J. C. Ho, “Robust irregular tensor factorization and completion for temporal health data analysis,” in *CIKM’20: Proc. 29th ACM Int. Conf. Inf. Knowl. Manag.*, 2020, pp. 1295–1304.
- [40] R. A. Harshman and M. E. Lundy, “Uniqueness proof for a family of models sharing features of tucker’s three-mode factor analysis and parafac/candecomp,” *Psychometrika*, vol. 61, no. 1, pp. 133–154, 1996.
- [41] J. M. ten Berge and H. A. Kiers, “Some uniqueness results for PARAFAC2,” *Psychometrika*, vol. 61, no. 1, pp. 123–132, 1996.
- [42] S. Friedland and G. Ottaviani, “The number of singular vector tuples and uniqueness of best rank-one approximation of tensors,” *Found. Comp. Math.*, vol. 14, no. 6, pp. 1209–1242, 2014.
- [43] R. Bro, R. Leardi, and L. G. Johnsen, “Solving the sign indeterminacy for multiway models,” *J. Chemom.*, vol. 27, no. 3-4, pp. 70–75, 2013.
- [44] N. E. Helwig, “The special sign indeterminacy of the direct-fitting Parafac2 model: Some implications, cautions, and recommendations for simultaneous component analysis,” *Psychometrika*, vol. 78, no. 4, pp. 725–739, 2013.
- [45] N. Gillis and F. Glineur, “Accelerated multiplicative updates and hierarchical ALS algorithms for nonnegative matrix factorization,” *Neural Comput.*, vol. 24, no. 4, pp. 1085–1105, 2012.
- [46] R. Bro and N. D. Sidiropoulos, “Least squares algorithms under unimodality and non-negativity constraints,” *J. Chemom.*, vol. 12, no. 4, pp. 223–247, 1998.
- [47] R. Bro, “Multi-way analysis in the food industry,” Ph.D. dissertation, Royal Veterinary and Agricultural University Denmark, 1998.
- [48] 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.
- [49] D. Bertsekas, *Nonlinear Programming*, ser. Athena Sci. Optim. Comput. Series. Athena Scientific, 1999, ISBN: 9781886529007.
- [50] N. Parikh and S. Boyd, “Proximal algorithms,” *Found. Trends Mach. Learn.*, vol. 1, no. 3, pp. 127–239, 2014.
- [51] A. Beck, *First-order methods in optimization*. SIAM, 2017.
- [52] L. Condat, “A direct algorithm for 1-D total variation denoising,” *IEEE Signal Process. Letters*, vol. 20, no. 11, pp. 1054–1057, 2013.
- [53] Q. F. Stout, “Unimodal regression via prefix isotonic regression,” *Comput. Stat. Data Anal.*, vol. 53, no. 2, pp. 289–297, 2008.
- [54] L. Condat, *Software*, [Accessed: 2020-10-19], 2019. [Online]. Available: <https://lcondat.github.io/software.html>.
- [55] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.

- [56] M. H. Van Benthem and M. R. Keenan, “Fast algorithm for the solution of large-scale non-negativity-constrained least squares problems,” *J. Chemom.*, vol. 18, no. 10, pp. 441–450, 2004.
- [57] J. Kim and H. Park, “Fast nonnegative matrix factorization: An active-set-like method and comparisons,” *SIAM J. Sci. Comput.*, vol. 33, no. 6, pp. 3261–3281, 2011.
- [58] B. J. Zijlstra and H. A. Kiers, “Degenerate solutions obtained from several variants of factor analysis,” *J. Chemom.*, vol. 16, no. 11, pp. 596–605, 2002.
- [59] R. L. Gollub *et al.*, “The MCIC collection: A shared repository of multi-modal, multi-site brain image data from a clinical investigation of schizophrenia,” *Neuroinformatics*, vol. 11, no. 3, pp. 367–388, 2013.
- [60] Z. Bai, P. Walker, A. Tschiffely, F. Wang, and I. Davidson, “Unsupervised network discovery for brain imaging data,” in *KDD’17: Proc. 23rd ACM SIGKDD Int. Conf. Knowl. Discov. and Data Mining*, 2017, pp. 55–64.
- [61] J. M. Amigo, T. Skov, R. Bro, J. Coello, and S. Maspoeh, “Solving GC-MS problems with PARAFAC2,” *TrAC Trends Anal. Chem.*, vol. 27, no. 8, pp. 714–725, 2008, ISSN: 0165-9936.
- [62] C. Chatzichristos, E. Kofidis, M. Morante, and S. Theodoridis, “Blind fMRI source unmixing via higher-order tensor decompositions,” *J. Neurosci. Methods*, vol. 315, pp. 17–47, 2019.
- [63] 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.
- [64] E. Gujral, R. Pasricha, and E. Papalexakis, “Beyond rank-1: Discovering rich community structure in multi-aspect graphs,” in *Proc. Web Conf. 2020*, 2020, pp. 452–462.
- [65] D. Mitchell, N. Ye, and H. De Sterck, “Nesterov acceleration of alternating least squares for canonical tensor decomposition: Momentum step size selection and restart mechanisms,” *Numer. Linear Algebra with Appl.*, vol. 27, no. 4, e2297, 2020.
- [66] A. M. Shun Ang, J. E. Cohen, L. T. Khanh Hien, and N. Gillis, “Extrapolated alternating algorithms for approximate canonical polyadic decomposition,” in *ICASSP’20: Proc. Int. Conf. Acoust., Speech, Signal Process.*, 2020, pp. 3147–3151.
- [67] H. Yu, D. Augustijn, and R. Bro, “Accelerating PARAFAC2 algorithms for non-negative complex tensor decomposition,” *Chemom. Intel. Lab. Syst.*, vol. 214, p. 104312, 2021.

An AO-ADMM approach to constraining PARAFAC2 on all modes*

Marie Roald[†] Carla Schenker[†] Vince D. Calhoun[‡]
Tülay Adalı[§] Rasmus Bro[¶] Jeremy E. Cohen^{||} Evrim Acar^{**}

1 Proofs

Proof of Proposition 3.1. Let $\mathcal{M} = \mathcal{O}_{J_1,R} \times \dots \times \mathcal{O}_{J_K,R} \times \mathbb{R}^{R \times R}$. There exists a continuous surjection between \mathcal{M} and \mathcal{P} [2], and it is, therefore, sufficient to prove the closedness of \mathcal{M} to prove Proposition 3.1. Closedness is preserved by cartesian products, and both $\mathbb{R}^{R \times R}$ and $\mathcal{O}_{J_k,R}$ are closed (closedness of $\mathcal{O}_{J_k,R}$ follows from the preimage description $\mathcal{O}_{J_k,R} = f^{-1}(\{\mathbf{I}\})$, where $f(\mathbf{X}) = \mathbf{X}^\top \mathbf{X}$). Thus, \mathcal{M} is closed, and consequently \mathcal{P} is also closed. \square

Proof of Theorem 4.1. Recall that $f : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}$ satisfies $f(a\mathbf{u}, a^{-2}\mathbf{v}) = f(\mathbf{u}, \mathbf{v})$, and that $r_u : \mathbb{R}^n \rightarrow \mathbb{R}$ and $r_v : \mathbb{R}^m \rightarrow \mathbb{R}$ are two (absolutely) homogeneous functions of degree d_u and d_v respectively. We want to show that

$$\min_{\mathbf{u}, \mathbf{v}} f(\mathbf{u}, \mathbf{v}) + ar_u(\mathbf{u}) + r_v(\mathbf{v}) \quad (1)$$

and

$$\min_{\mathbf{u}, \mathbf{v}} f(\mathbf{u}, \mathbf{v}) + r_u(\mathbf{u}) + a^{2\frac{d_u}{d_v}} r_v(\mathbf{v}) \quad (2)$$

are equivalent.

We start by introducing the change of variables $\mathbf{u} = a^{-\frac{1}{d_u}} \tilde{\mathbf{u}}$ and $\mathbf{v} = a^{\frac{2}{d_u}} \tilde{\mathbf{v}}$ into Equation (1):

$$\min_{\tilde{\mathbf{u}}, \tilde{\mathbf{v}}} f(a^{-\frac{1}{d_u}} \tilde{\mathbf{u}}, a^{\frac{2}{d_u}} \tilde{\mathbf{v}}) + ar_u(a^{-\frac{1}{d_u}} \tilde{\mathbf{u}}) + r_v(a^{\frac{2}{d_u}} \tilde{\mathbf{v}}). \quad (3)$$

Next, from $f(a\mathbf{u}, a^{-2}\mathbf{v}) = f(\mathbf{u}, \mathbf{v})$ and the homogeneity of r_u and r_v , we have

$$\min_{\tilde{\mathbf{u}}, \tilde{\mathbf{v}}} f(\tilde{\mathbf{u}}, \tilde{\mathbf{v}}) + a \left(a^{-\frac{1}{d_u}} \right)^{d_u} r_u(\tilde{\mathbf{u}}) + \left(a^{\frac{2}{d_u}} \right)^{d_v} r_v(\tilde{\mathbf{v}}), \quad (4)$$

***Funding:** This work was supported in part by the Research Council of Norway through project 300489 (IKTPLUS) and ANR JCJC project LoRAiA ANR-20-CE23-0010.

[†]Department of Data Science and Knowledge Discovery, Simula Metropolitan Center for Digital Engineering, Oslo, Norway & Faculty of Technology, Art and Design, Oslo Metropolitan University, Oslo, Norway (mariero@simula.no, carla@simula.no).

[‡]Department of Psychology, Georgia State University, Atlanta, GA, USA (vcalhoun@gsu.edu)

[§]Department of Computer Science and Electrical Engineering, UMBC, Baltimore, MD (adali@umbc.edu).

[¶]Department of Food Science, University of Copenhagen, Copenhagen, Denmark (rb@food.ku.dk).

^{||}Univ Lyon, INSA-Lyon, UCBL, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1206, F-69100 Villeurbanne, France (jeremy.cohen@cnrs.fr).

**Department of Data Science and Knowledge Discovery, Simula Metropolitan Center for Digital Engineering, Oslo, Norway (evrim@simula.no).

which is equivalent to

$$\min_{\tilde{\mathbf{u}}, \tilde{\mathbf{v}}} f(\tilde{\mathbf{u}}, a\tilde{\mathbf{v}}) + r_u(\tilde{\mathbf{u}}) + a^{2\frac{d_v}{d_u}} r_v(\tilde{\mathbf{v}}). \quad (5)$$

□

2 Additional observations on \mathcal{P}

2.1 Computations for the one-component case

In the one-component case, we have $\mathcal{O}_{J,1} = \{\mathbf{x} \in \mathbb{R}^J \mid \|\mathbf{x}\| = 1\}$ and $\mathbf{B}_k^\top \mathbf{B}_k = \delta^2 \in \mathbb{R}$. Consequently, $\delta\mathcal{O}_{J,1}$ is a 2-norm sphere with radius δ . Solving

$$\min_{\delta} \sum_{k=1}^K d(\delta\mathcal{O}_{J,1}, \mathbf{B}_k) = \min_{\delta} \sum_{k=1}^K \min_{\mathbf{P}_k \in \mathcal{O}_{J,1}} \|\mathbf{B}_k - \mathbf{P}_k \delta\|_F^2 \quad (6)$$

is therefore a bilevel optimization problem where the inner problem is the projection onto the sphere with radius δ , which has the closed form solution (for $\mathbf{B}_k \neq 0$):

$$\arg \min_{\mathbf{P}_k \in \mathcal{O}_{J,1}} \|\mathbf{B}_k - \mathbf{P}_k \delta\|_F^2 = \frac{[\mathbf{b}_k]_1}{\|[\mathbf{b}_k]_1\|}. \quad (7)$$

Substituting Equation (7) into Equation (6) yields

$$\min_{\delta} \sum_{k=1}^K d(\delta\mathcal{O}_{J,1}, \mathbf{B}_k) = \min_{\delta} \sum_{k=1}^K \left\| \mathbf{B}_k - \frac{[\mathbf{b}_k]_1}{\|[\mathbf{b}_k]_1\|} \delta \right\|_F^2, \quad (8)$$

which is a formulation of the mean. For the one-component case, we therefore have that projection onto \mathcal{P} is equivalent to projecting onto the sphere with radius given by the average norm of each \mathbf{B}_k .

2.2 Additional interpretations of \mathcal{P}

One interpretation of \mathcal{P} arises from the formulation discussed in Subsection 3.1. We can interpret each \mathbf{P}_k matrix as an orthogonal basis for a given subspace and $\Delta_{\mathbf{B}}$ as a matrix of coordinates for each component in that subspace.

Another interpretation of \mathcal{P} , for the case where $J_k = J$, is apparent by observing that

$$\mathbf{B}_k = \mathbf{P}_k \Delta_{\mathbf{B}} = \mathbf{P}_k \mathbf{P}_1^\top \mathbf{P}_1 \Delta_{\mathbf{B}} = \tilde{\mathbf{P}}_k \mathbf{B}_1, \quad (9)$$

where $\tilde{\mathbf{P}}_k = \mathbf{P}_k \mathbf{P}_1^\top$. Thus, every \mathbf{B}_k matrix can be constructed by rotating and reflecting the columns of \mathbf{B}_1 . See Figure 1 for an illustration of this with two slabs constructed from two components in \mathbb{R}^2 .

3 Some proximable constraints

3.1 Hard constraints

Imposing hard constraints, that is, restricting the component vectors to lie in a specific set, requires evaluating the projection onto that set. For some constraints, the projection has a fast closed-form solution (e.g., non-negativity). Other constraints

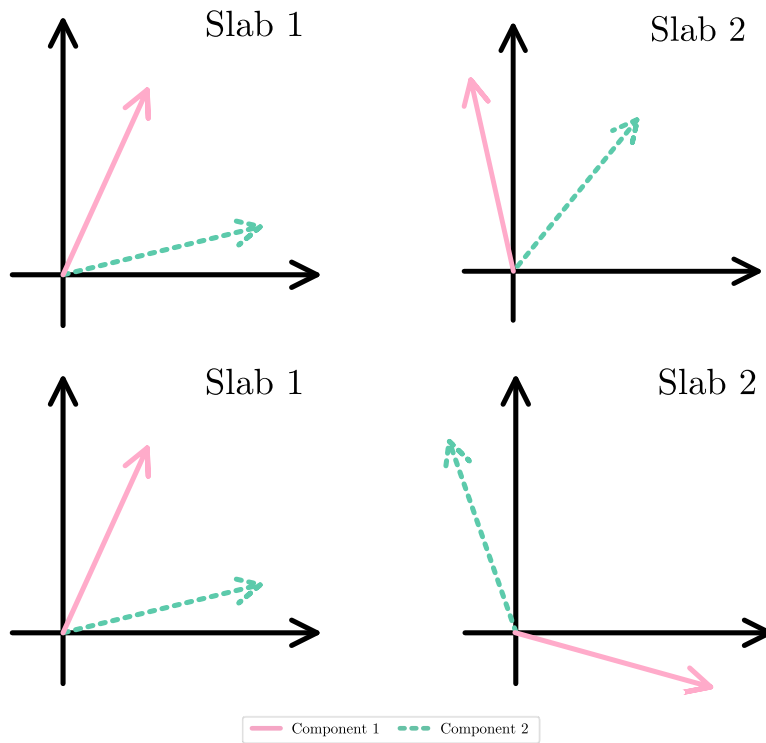


Figure 1: Illustration of the PARAFAC2 constraint with two components in two dimensions across two slabs. Each row represents one set of evolving component vectors and each column represents tensor slabs. The component vectors in the top row satisfy the PARAFAC2 constraint, since the angle between the components are constant for the first and second slab. The component vectors in the bottom row do not satisfy the PARAFAC2 constraint, since the angle between them are not constant for the two slabs.

are more involved and may not even be convex. For example, unimodality constraints require that we solve a set of isotonic regression problems for each component vector, and the proximal operator is not unique [3], [4]. The PARAFAC2 constraint is another example of a non-convex hard constraint.

3.2 Graph Laplacian regularization

For graph Laplacian regularization, we penalize pairwise differences between component vector elements. With appropriately selected weight parameters $w_{ij} \geq 0$, we get the graph Laplacian penalty

$$g(\mathbf{x}) = \sum_{i,j} w_{ij} (x_i - x_j)^2 = \mathbf{x}^\top \mathbf{L} \mathbf{x}, \quad (10)$$

where \mathbf{L} is a positive semidefinite matrix with $l_{ij} = -w_{ij}$ for $i \neq j$ and $l_{ii} = \sum_j w_{ij}$. The scaled proximal operator of this penalty function is given by

$$\text{prox}_{\frac{\mathbf{x}^\top \mathbf{L} \mathbf{x}}{\rho}}(\mathbf{x}) = \left(\mathbf{L} + \frac{\rho}{2} \mathbf{I} \right)^{-1} \frac{\rho}{2} \mathbf{x}, \quad (11)$$

which can be evaluated using the Cholesky factorization or algebraic multigrid preconditioned conjugate gradient [5].

3.3 TV regularization

Another popular regularization technique is TV regularization, which encourages piecewise constant components and is defined as

$$g(\mathbf{x}) = \sum_i |x_{i+1} - x_i| \quad (12)$$

While there is no closed form solution for the proximal operator of the TV penalty, we can still evaluate it in linear time [6].

4 Including ridge regularization in the data fidelity updates

Here, we state the update steps with ridge imposed for all modes. For the \mathbf{A} -updates, the update becomes

$$\text{prox}_{\frac{f_{\mathbf{A}+\gamma} \|\cdot\|_F^2}{\rho_{\mathbf{A}}}}(\mathbf{M}) = \left(\sum_{k=1}^K \mathbf{X}_k \mathbf{\Gamma}_k + \frac{\rho_{\mathbf{A}}}{2} \mathbf{M} \right) \left(\sum_{k=1}^K \mathbf{\Gamma}_k^\top \mathbf{\Gamma}_k + \frac{2\gamma + \rho_{\mathbf{A}}}{2} \mathbf{I} \right)^{-1}, \quad (13)$$

if we impose ridge regularization with a strength of γ . Likewise, the updates with respect to \mathbf{B}_k and \mathbf{D}_k become

$$\mathbf{B}_k \leftarrow \left(\mathbf{X}_k^\top \mathbf{A} \mathbf{D}_k + \frac{\rho_{\mathbf{B}_k}}{2} \mathbf{M} \right) \left(\mathbf{D}_k \mathbf{A}^\top \mathbf{A} \mathbf{D}_k + (\gamma + \rho_{\mathbf{B}_k}) \mathbf{I} \right)^{-1}, \quad (14)$$

where $\mathbf{M} = (\mathbf{Z}_{\mathbf{B}_k} - \boldsymbol{\mu}_{\mathbf{Z}_{\mathbf{B}_k}} + \mathbf{Y}_{\mathbf{B}_k} - \boldsymbol{\mu}_{\Delta_{\mathbf{B}_k}})$, and

$$\text{prox}_{\frac{f_{\mathbf{D}_k+\gamma} \|\cdot\|_F^2}{\rho_{\mathbf{D}_k}}}(\mathbf{v}) = \left(\mathbf{A}^\top \mathbf{A} * \mathbf{B}_k^\top \mathbf{B}_k + \frac{2\gamma + \rho_{\mathbf{D}_k}}{2} \mathbf{I} \right)^{-1} \left(\text{Diag}(\mathbf{A}^\top \mathbf{X}_k \mathbf{B}_k) + \frac{\rho_{\mathbf{D}_k}}{2} \mathbf{v} \right), \quad (15)$$

respectively.

5 CP-based updates for the non-evolving modes

To obtain the CP-based update steps, we use the reformulation in [2]. We start by assuming that the \mathbf{B}_k factors are feasible, which allows the substitution $\mathbf{B}_k = \mathbf{Y}_{\mathbf{B}_k}$. Setting $\mathbf{T}_k = \mathbf{X}_k \mathbf{P}_k$, we see that updating \mathbf{A} is equivalent to solving

$$\min_{\mathbf{A}} \sum_{k=1}^K \left\| \mathbf{A} \mathbf{D}_k \mathbf{\Delta}_k \mathbf{\Delta}_k^\top - \mathbf{T}_k \right\|_F + g_{\mathbf{A}}(\mathbf{A}), \quad (16)$$

which is the \mathbf{A} -updates of a regularized CP-model fitted to \mathcal{J} (i.e. the tensor with frontal slices given by \mathbf{T}_k). We can use a similar argument for $\{\mathbf{D}_k\}_{k \leq K}$, yielding the update rules from [7], [8]:

$$\text{prox}_{\frac{f_{\mathbf{A}}}{\rho_{\mathbf{A}}}}(\mathbf{M}) = \left(\mathbf{T}_{(1)} (\mathbf{\Delta}_{\mathbf{B}} \odot \mathbf{C}) + \frac{\rho_{\mathbf{A}}}{2} \mathbf{M} \right) \left(\mathbf{\Delta}_{\mathbf{B}}^\top \mathbf{\Delta}_{\mathbf{B}} * \mathbf{C}^\top \mathbf{C} + \frac{\rho_{\mathbf{A}}}{2} \mathbf{I} \right)^{-1} \quad (17)$$

and

$$\text{prox}_{\frac{f_{\mathbf{C}}}{\rho_{\mathbf{C}}}}(\mathbf{M}) = \left(\mathbf{T}_{(3)} (\mathbf{A} \odot \mathbf{\Delta}_{\mathbf{B}}) + \frac{\rho_{\mathbf{C}}}{2} \mathbf{M} \right) \left(\mathbf{A}^\top \mathbf{A} * \mathbf{\Delta}_{\mathbf{B}}^\top \mathbf{\Delta}_{\mathbf{B}} + \frac{\rho_{\mathbf{C}}}{2} \mathbf{I} \right)^{-1}, \quad (18)$$

where $\mathbf{T}_{(n)}$ is the mode- n unfolding of \mathcal{J} .

For both the CP- and the CMF-based updates, we get the same overall ADMM algorithm for \mathbf{A} ; the only difference is which linear system we need to solve. However, for the updates of the \mathbf{D}_k -matrices, we see clear differences. For the CMF-based updates, we solve a separate system of equations for each \mathbf{D}_k -matrix (i.e. each row of \mathbf{C}), while for the CP-based updates, all \mathbf{D}_k -matrices are obtained by solving the same system of equations.

To select the penalty parameters with the CP-based updates, we used the following heuristic:

$$\rho_{\mathbf{A}} = \frac{1}{R} \text{Tr} \left(\mathbf{\Delta}_{\mathbf{B}}^\top \mathbf{\Delta}_{\mathbf{B}} * \mathbf{C}^\top \mathbf{C} \right), \quad \rho_{\mathbf{C}} = \frac{1}{R} \text{Tr} \left(\mathbf{A}^\top \mathbf{A} * \mathbf{\Delta}_{\mathbf{B}}^\top \mathbf{\Delta}_{\mathbf{B}} \right). \quad (19)$$

Algorithm 1: CMF-based and CP-based ADMM for the \mathbf{A} -matrices

Result: $\mathbf{A}, \mathbf{Z}_{\mathbf{A}}, \mu_{\mathbf{A}}$

- 1 **while** stopping conditions are not met and max no. iterations not exceeded **do**
- 2 $\mathbf{A} \leftarrow \xrightarrow{\text{Equation (17)}} \text{prox}_{\frac{f_{\mathbf{A}}}{\rho_{\mathbf{A}}}}(\mathbf{Z}_{\mathbf{A}} - \mu_{\mathbf{A}})$
- 3 $\mathbf{Z}_{\mathbf{A}} \leftarrow \text{prox}_{\frac{g_{\mathbf{A}}}{\rho_{\mathbf{A}}}}(\mathbf{A} + \mu_{\mathbf{A}})$
- 4 $\mu_{\mathbf{A}} \leftarrow \mu_{\mathbf{A}} + \mathbf{A} - \mathbf{Z}_{\mathbf{A}}$
- 5 **end**

5.1 Computational complexity

One difference between the CP-based and CMF-based update steps is their computational complexity, which are given in Table 1.

From the complexities alone, it seems like the CP-based updates have lower computational complexity. However, for the CP-based updates, we also need to compute the \mathcal{J} tensor, by multiplying the frontal slices of \mathcal{X} with the \mathbf{P}_k -matrices, which incurs an additional $O(IJKR)$ step in the algorithm. Also, the \mathbf{B}_k updates have a substantial time-complexity that may dominate the \mathbf{A} and \mathbf{D}_k updates. Finally, with the

Algorithm 2: CP-based ADMM for the \mathbf{D}_k -matrices (\mathbf{C} -matrix)

Result: $\mathbf{C}, \mathbf{Z}_\mathbf{C}, \boldsymbol{\mu}_\mathbf{C}$

```

1 while stopping conditions are not met and max no. iterations not exceeded do
2    $\mathbf{C} \leftarrow \text{Equation (18)} \text{prox}_{\frac{f_\mathbf{C}}{\rho_\mathbf{C}}}(\mathbf{Z}_\mathbf{C} - \boldsymbol{\mu}_\mathbf{C})$ 
3    $\mathbf{Z}_\mathbf{C} \leftarrow \text{prox}_{\frac{g_\mathbf{C}}{\rho_\mathbf{C}}}(\mathbf{C} + \boldsymbol{\mu}_\mathbf{C})$ 
4    $\boldsymbol{\mu}_\mathbf{C} \leftarrow \boldsymbol{\mu}_\mathbf{C} + \mathbf{C} - \mathbf{Z}_\mathbf{C}$ 
5 end

```

Table 1: Computational complexities for the different AO-ADMM update steps. I , J and K denote the tensor size, R denotes the number of components and Q the number of iterations.

	CMF-based updates	CP-based updates
\mathbf{A}	$O(IJKR + JKR^2 + R^3 + IR^2Q)$	$O(IKR^2 + R^3 + IR^2Q)^*$
$\{\mathbf{D}_k\}_{k \leq K}$	$O(IJKR + IR^2 + JKR^2 + KR^3 + KR^2Q)$	$O(IKR^2 + R^3 + KR^2Q)^*$
$\{\mathbf{B}_k\}_{k \leq K}$	$O(IJKR + IR^2 + KR^3 + JKR^2Q)$	—

* The CP-based update steps also require computing the \mathcal{J} -tensor which adds an additional $O(IJKR)$ complexity step shared between the A and D updates.

CMF-based updates, we update the \mathbf{A} - and \mathbf{D}_k -matrices based on the \mathbf{B}_k -matrices. However, for the CP-based updates, we update these matrices based on the $\mathbf{Y}_{\mathbf{B}_k}$ -matrices. Thus, the two algorithms may behave differently, and it is difficult to select the “optimal” update steps based on computational complexity alone.

5.2 Setup SM1: Comparing the CMF- and CP-based ADMM updates

Dataset generation To compare the two ADMM update schemes for \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$, we generated non-negative \mathbf{B}_k -factor matrices by first creating a non-negative $\boldsymbol{\Delta}_\mathbf{B}$ and non-negative $\{\mathbf{P}_k\}_{k \leq K}$ and then setting $\mathbf{B}_k = \mathbf{P}_k \boldsymbol{\Delta}_\mathbf{B}$. The elements of $\boldsymbol{\Delta}_\mathbf{B}$ were drawn from a uniform distribution between 0 and 1. Each \mathbf{P}_k -matrix was generated by first dividing the indices between 1 and J into R contiguous partitions, each containing the non-zero elements of one column of \mathbf{P}_k , and then drawing these non-zero elements from a uniform distribution between 0 and 1. The partition indices were generated from an R -category Dirichlet distribution (with each concentration parameter set to 1) scaled by J , rounded to the nearest integer. We used this approach to generate 30 five-component PARAFAC2 decompositions of various sizes. Ten of size $30 \times 10 \times 70$, ten of size $30 \times 100 \times 70$ and ten of size $30 \times 1000 \times 70$. For each of these decompositions, we generated two simulated data tensors with different noise levels; one with $\eta = 0.33$ and one with $\eta = 0.5$.

Experiment setup Each data tensor was decomposed with ten random initialization using both the CMF-based updates and the CP-based updates. Each initialization ran until convergence or for at most 2000 outer iterations. With both schemes, we imposed non-negativity on all modes.

Table 2: Mean performance (\pm one standard deviation) for the selected initialization for all simulated datasets with the CMF-based and CP-based AO-ADMM scheme.

η	J	FMS		Final iteration		Time [s]	
		CMF	CP	CMF	CP	CMF	CP
0.33	10^1	0.98 ± 0	0.98 ± 0	195 ± 77	192 ± 64	11 ± 5	10 ± 3
	10^2	0.98 ± 0	0.98 ± 0	179 ± 81	188 ± 116	50 ± 24	46 ± 30
	10^3	0.98 ± 0	0.98 ± 0	144 ± 65	160 ± 71	98 ± 44	104 ± 46
0.5	10^1	0.96 ± 0.01	0.96 ± 0.01	301 ± 289	320 ± 389	14 ± 9	13 ± 11
	10^2	0.96 ± 0.01	0.96 ± 0.01	150 ± 56	151 ± 37	41 ± 16	33 ± 8
	10^3	0.96 ± 0.01	0.96 ± 0.01	141 ± 56	141 ± 52	93 ± 35	86 ± 32

Table 3: Setup SM1: Mean performance (\pm one standard deviation) across all initializations and all simulated data tensors with the CMF-based and CP-based AO-ADMM scheme.

η	J	FMS		Final iteration		Time [s]	
		CMF	CP	CMF	CP	CMF	CP
0.33	10^1	0.98 ± 0	0.98 ± 0	261 ± 302	342 ± 461	15 ± 18	18 ± 24
	10^2	0.98 ± 0	0.98 ± 0	179 ± 142	227 ± 280	51 ± 46	59 ± 83
	10^3	0.98 ± 0	0.98 ± 0	152 ± 65	167 ± 149	105 ± 43	109 ± 106
0.5	10^1	0.96 ± 0.01	0.96 ± 0.01	525 ± 647	699 ± 777	27 ± 35	33 ± 39
	10^2	0.96 ± 0.01	0.96 ± 0.01	156 ± 88	182 ± 185	43 ± 29	47 ± 60
	10^3	0.96 ± 0.01	0.96 ± 0.01	140 ± 60	142 ± 62	93 ± 41	87 ± 38

Results Tables 2 and 3 and Figures 2 and 3 demonstrate that the two AO-ADMM schemes are comparable in terms of time and FMS.

6 Additional details for the simulation experiments

6.1 Setup 2

To generate \mathbf{B}_k -factor matrices, solved the optimization problem

$$\begin{aligned} \min_{\mathbf{B}_k} \quad & \|\mathbf{B}_k^\top \mathbf{B}_k - \mathbf{X}^\top \mathbf{X}\|^2 \\ \text{subject to} \quad & [\mathbf{B}_k]_{jr} \geq 0 \end{aligned} \tag{20}$$

using projected gradient descent with different standard normal random initializations for each \mathbf{B}_k -factor matrix and a fixed \mathbf{X} matrix with elements from a truncated standard normal distribution. Then, if projected gradient descent did not obtain a loss-value less than 10^{-12} within 10000 iterations, we re-initialized \mathbf{B}_k , reduced the step-size by a factor 10 and tried again. If no suitable \mathbf{B}_k factor matrix could be found after 10 attempts, we selected a new \mathbf{X} -matrix and restarted the process.

Figure 4 shows the results for noise levels $\eta = 0.6, 0.85, 1.2, 1.7$ and 2.5 , and Table 4 gives an overview of the number of datasets left after removing datasets where ALS gave only degenerate solutions.

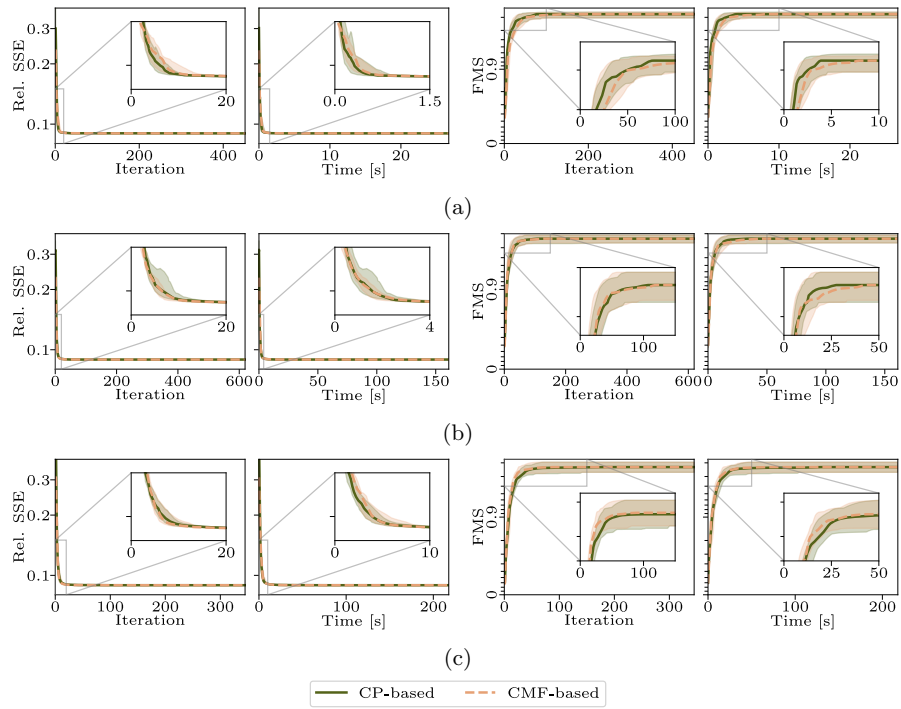


Figure 2: Setup SM1: Performance plots on the data tensors with $\eta = 0.33$ for the CMF-based and CP-based AO-ADMM schemes. $J=10, 100,$ and 1000 in (a), (b) and (c), respectively

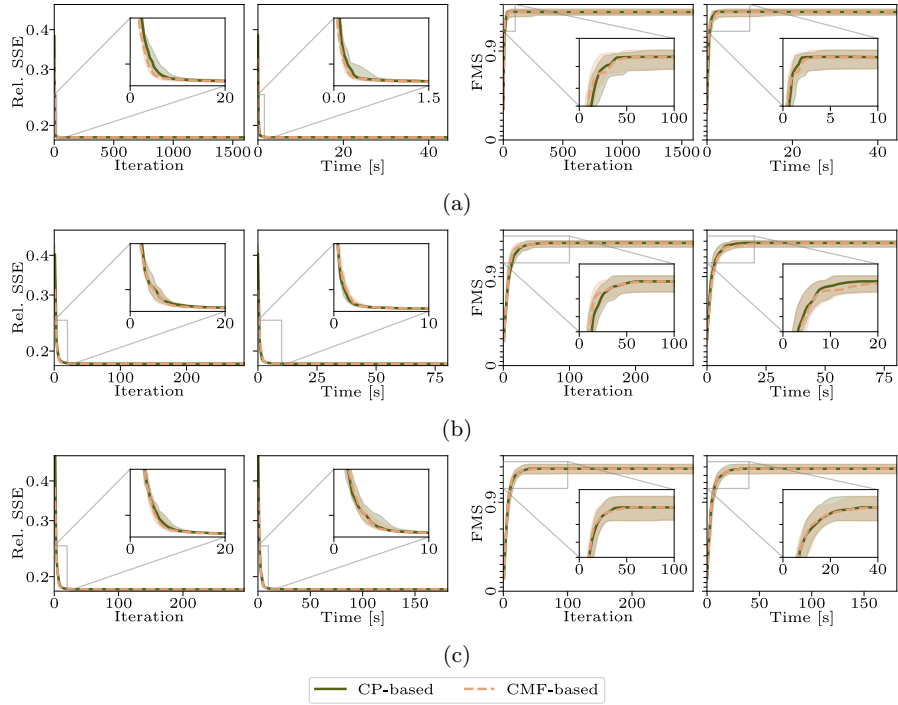


Figure 3: Setup SM1: Performance plots on the data tensors with $\eta = 0.5$ for the CMF-based and CP-based AO-ADMM schemes. $J=10, 100$, and 1000 in (a), (b) and (c), respectively

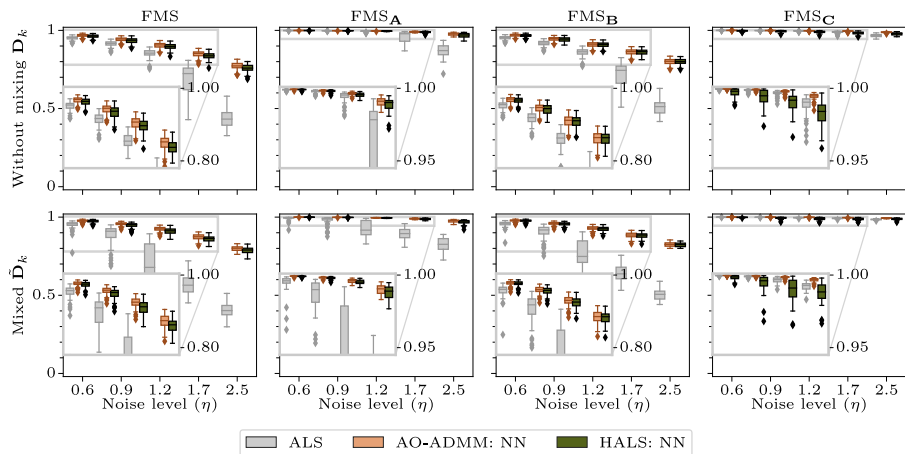


Figure 4: Setup 2: Boxplots showing the FMS for the different datasets.

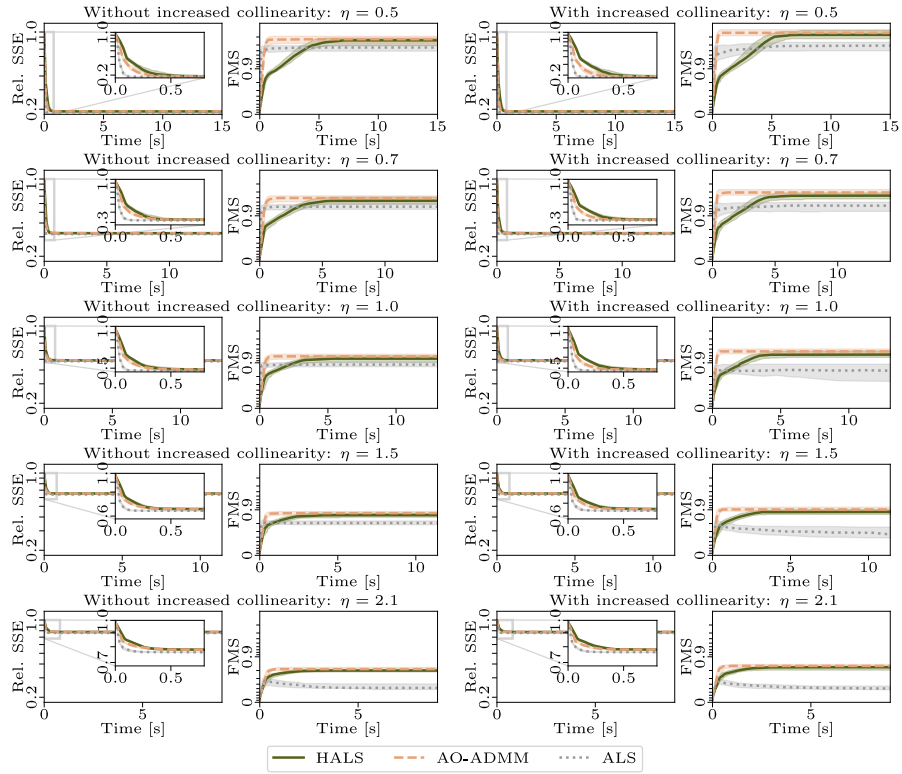


Figure 5: Setup 2: Performance plots showing FMS and relative SSE plotted against time. Each row corresponds to a different noise level.




Figure 6: Setup 2: Performance plots showing FMS and relative SSE plotted against iteration. Each row corresponds to a different noise level.

Table 4: Setup 2: Number of datasets where ALS gave at least one non-degenerate solution (of a total of 50 datasets)

Noise level	0.5	0.6	0.7	0.9	1.0	1.2	1.5	1.7	2.1	2.5
Data										
No mixing	50	50	50	50	50	50	50	49	49	39
Mixing C	50	50	50	50	50	50	49	39	32	23




Figure 7: Setup 3: Performance plots showing FMS and relative SSE plotted against time and iteration.

6.2 Setup 3

In Figure 7, we see that the models fitted with AO-ADMM with only non-negativity constraints converged faster than those fitted with the HALS scheme. Furthermore, we observe that while the models fitted with unimodality constraint took longer to converge than the other models, they obtained a much higher FMS.

6.3 Setup 3b

Data generation We also ran experiments using the same data generation setup as for Setup 3, but with a constant $\sigma_{kr} = \sigma_r$ for all k , which means that the only change is a constant shifting for all components. Thus, the \mathbf{B}_k -components followed the PARAFAC2 constraint.

Experiment settings We used the same experimental settings as for Setup 3 to decompose these datasets (including the initialization scheme and the scaling of $\rho_{\mathbf{B}_k}$ and the extra ADMM iterations).

Results In Figure 8, we see the results from these experiments. For six datasets, ALS yielded degenerate solutions for all 20 initializations, leaving 44 out of 50 datasets. We see from the figures that the results are similar to those from Setup 3. Specifically, we see that imposing constraints on the evolving mode improves factor recovery, especially when the component vectors are constrained to be unimodal. Figure 9 shows the true and estimated components for one of the datasets.

6.4 Setup 4 and 5

In Tables 5 and 6, we see the number of initializations that converged and the number of simulated datasets where at least one initialization converged for the models fitted with graph Laplacian regularization and TV regularization, respectively. Also, in Figures 10 and 11, we see the regularization penalty (after normalizing the \mathbf{B}_k -matrices)




Figure 8: Setup 3: Boxplots showing the FMS for different models fitted to datasets with unimodal \mathbf{B}_k -matrices. NN represents models fitted with non-negativity imposed on all modes using AO-ADMM, NN&U represents models fitted with non-negativity on all modes and unimodality imposed on $\{\mathbf{B}_k\}_{k \leq K}$ using AO-ADMM, and ALS represents models fitted with non-negativity imposed on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$.

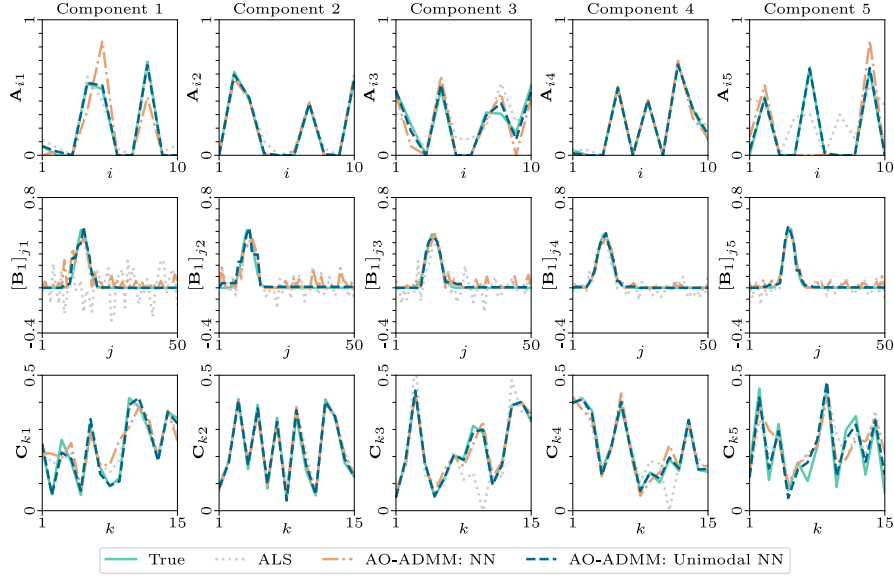


Figure 9: Setup 3: Plots showing the true and estimated components for one of the datasets. NN represents the model fitted with non-negativity imposed on all modes, Unimodal NN represents the model fitted with non-negativity on all modes and unimodality imposed on $\{\mathbf{B}_k\}_{k \leq K}$ and ALS represents the model fitted with non-negativity imposed on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. The cwSNRs (left to right) are $-2.4, -9.3, -12, 4.5$ and -11 dB for this frontal slice.

Table 5: Setup 4 and 5: Number of converged initializations and datasets with at least one converged initialization for the models fitted with graph Laplacian regularization, the *Successful decompositions* column shows the number of datasets where at least one initialization converged. Convergence for ALS is measured in the same way using the same tolerance as AO-ADMM (with no feasibility gaps due to ALS not being a splitting method).

Method	Smooth	Converged init.	Successful decompositions
ADMM - Ridge=0.0	1	0/400	0/20
	10	0/400	0/20
	100	0/400	0/20
	1000	0/400	0/20
ADMM - Ridge=0.1	1	272/400	19/20
	10	255/400	18/20
	100	261/400	18/20
	1000	261/400	17/20
ALS	N/A	393/400	20/20

as a function of iteration number for the selected initialization for each simulated dataset. Once an algorithm has converged, its final value is set as constant for the rest of the iterations, making it easier to compare the behavior of the ALS algorithm and AO-ADMM algorithm without ridge penalty on \mathbf{A} and $\{\mathbf{D}_k\}_{k \leq K}$. Finally, in Figure 12, we see a boxplot depicting the FMS of the model fitted with graph Laplacian regularization and TV regularization when we select the initialization that obtained the highest overall FMS.

Table 6: Setup 4 and 5: Number of converged initializations and datasets with at least one converged initialization for the models fitted with TV regularization, the *Successful decompositions* column shows the number of datasets where at least one initialization converged. Convergence for ALS is measured in the same way using the same tolerance as AO-ADMM (with no feasibility gaps due to ALS not being a splitting method).

Method	Reg	Converged init.	Successful decompositions
ADMM - Ridge=0.0	0.001	0/400	0/20
	0.01	0/400	0/20
	0.1	0/400	0/20
	1	0/400	0/20
	10	0/400	0/20
ADMM - Ridge=0.1	0.001	38/400	7/20
	0.01	198/400	16/20
	0.1	230/400	20/20
	1	199/400	18/20
	10	54/400	8/20
ALS	N/A	398/400	20/20

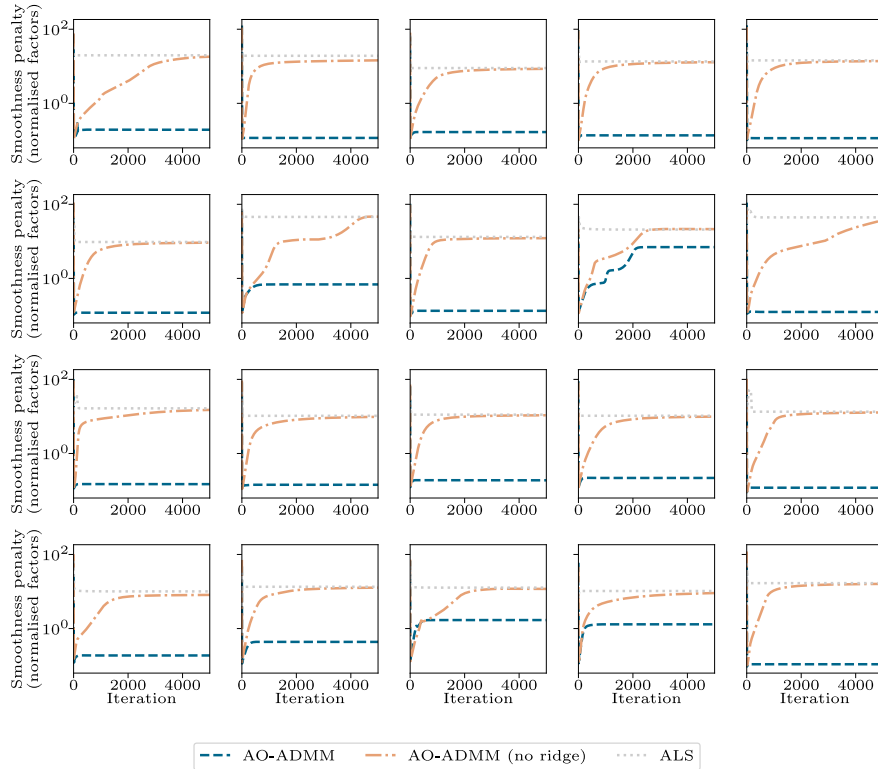


Figure 10: Setup 4: Graph Laplacian penalty of components (after normalizing) as a function of iteration number. Each plot corresponds to one dataset.

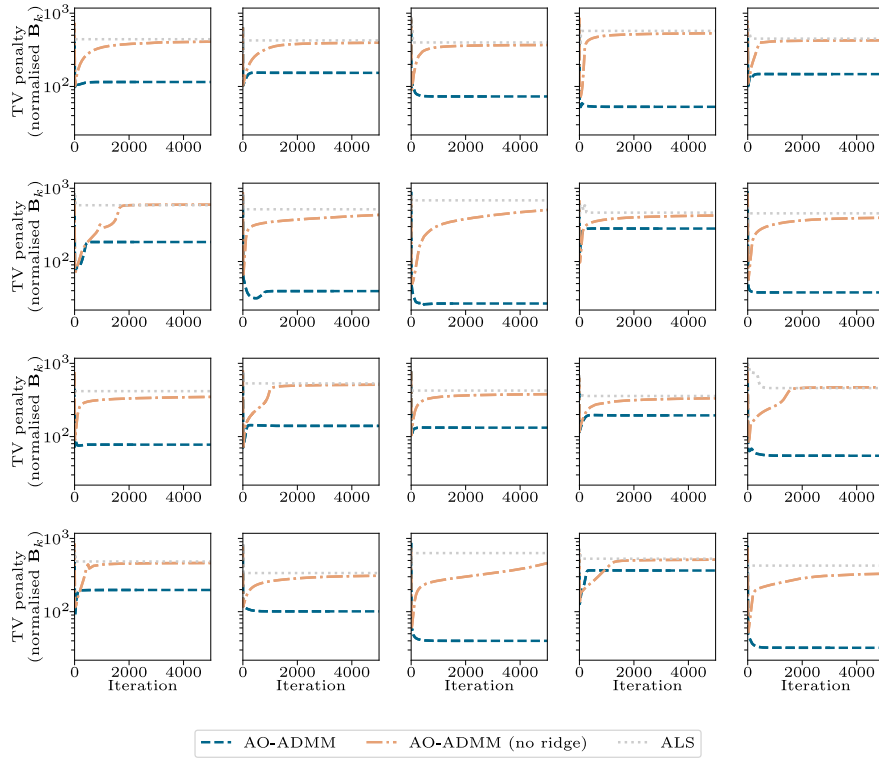


Figure 11: Setup 5: Total variation of components (after normalizing) as a function of iteration number. Each plot corresponds to one dataset.

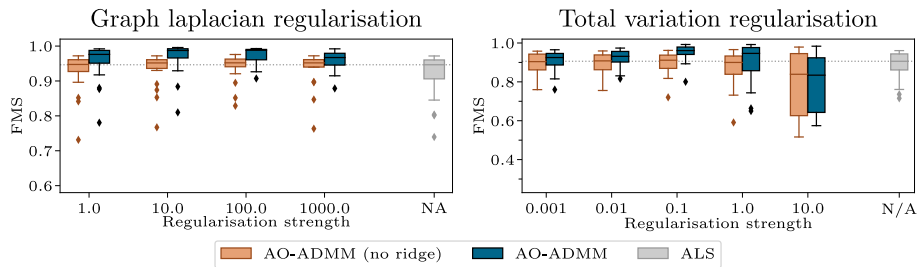


Figure 12: Setup 4 and 5: The FMS for the different models fitted with penalty based regularization when the initialization that obtained the highest overall FMS is chosen for each simulated dataset.

References

- [1] M. Roald, C. Schenker, J. E. Cohen, and E. Acar, “PARAFAC2 AO-ADMM: Constraints in all modes,” in *EUSIPCO’21: Proc. 2021 29th Eur. Signal Process. Conf.*, EURASIP, 2021.
- [2] H. A. L. Kiers, J. M. F. Ten Berge, and R. Bro, “PARAFAC2 - Part I. A direct fitting algorithm for the PARAFAC2 model,” *J. Chemom.*, vol. 13, no. 3-4, pp. 275–294, 1999.
- [3] R. Bro and N. D. Sidiropoulos, “Least squares algorithms under unimodality and non-negativity constraints,” *J. Chemom.*, vol. 12, no. 4, pp. 223–247, 1998.
- [4] Q. F. Stout, “Unimodal regression via prefix isotonic regression,” *Comput. Stat. & Data Anal.*, vol. 53, no. 2, pp. 289–297, 2008.
- [5] O. E. Livne and A. Brandt, “Lean algebraic multigrid (lamg): Fast graph laplacian linear solver,” *SIAM J. Sci. Comput.*, vol. 34, no. 4, B499–B522, 2012.
- [6] L. Condat, “A direct algorithm for 1-D total variation denoising,” *IEEE Signal Process. Letters*, vol. 20, no. 11, pp. 1054–1057, 2013.
- [7] K. Huang, N. D. Sidiropoulos, and A. P. Liavas, “A flexible and efficient algorithmic framework for constrained matrix and tensor factorization,” *IEEE Trans. Signal Process.*, vol. 64, no. 19, pp. 5052–5065, 2016.
- [8] A. Afshar, I. Perros, E. E. Papalexakis, E. Searles, J. Ho, and J. Sun, “COPA: Constrained PARAFAC2 for Sparse & Large Datasets,” in *ACM Int. Conf. on Inf. and Knowl. Management*, 2018, pp. 793–802.