

A SMOOTHED AND PROBABILISTIC PARAFAC MODEL WITH COVARIATES

BY LEYING GUAN¹

¹*Department of Biostatistics, Yale University, leying.guan@yale.edu*

Analysis and clustering of multivariate time-series data attract growing interest in immunological and clinical studies. In such applications, researchers are interested in clustering subjects based on potentially high-dimensional longitudinal features, and in investigating how clinical covariates may affect the clustering results. These studies are often challenging due to high dimensionality, as well as the sparse and irregular nature of sample collection along the time dimension. We propose a smoothed probabilistic PARAFAC model with covariates (SPACO) to tackle these two problems while utilizing auxiliary covariates of interest. We provide intensive simulations to test different aspects of SPACO and demonstrate its use on immunological data sets from two recent cohorts of SARs-CoV-2 patients.

1. Introduction. We consider the modeling of sparsely observed multivariate longitudinal data. Such data are common in medical applications, where we may observe for each participant/subject $i = 1, \dots, I$ a matrix $\mathbf{X}_i \in \mathbb{R}^{n_i \times J}$, measured for J different features and collected at n_i different times $\{t_{i,1}, \dots, t_{i,n_i}\}$. For example, immune profiles are measured for hundreds of markers for each patient at irregular sampling times in [Lucas et al. \(2020\)](#) and [Rendeiro et al. \(2020\)](#). Letting T be the number of distinct time points across all subjects, such multivariate time-series data can be viewed as a three-way tensor $\mathbf{X} \in \mathbb{R}^{I \times T \times J}$, with potentially high missing-rate along the time dimension.

In applications such as immunological analysis, \mathbf{X} usually has many features J , with strong correlation structure. Moreover, the data for each subject i tends to be sparsely and irregularly sampled in time. In addition to \mathbf{X} , researchers often have a set of nontemporal covariates $\mathbf{z}_i \in \mathbb{R}^q$ associated with subject i . These auxiliary covariates are perhaps helpful in accounting for the heterogeneity in $\mathbf{X}_{ij}(t)$ across different subjects. Researchers are often interested in understanding how subjects differ in terms of features' time trajectories, and how these differences relate to these covariates $\mathbf{Z} := (\mathbf{z}_1, \dots, \mathbf{z}_I)^\top$.

To adapt to the characteristics of \mathbf{X} and efficiently utilize the auxiliary variables \mathbf{Z} , we propose a smoothed and probabilistic PARAFAC model for \mathbf{X} in the presence of auxiliary covariates (SPACO). SPACO assumes that (1) \mathbf{X} is a noisy realization of an underlying low-rank tensor, and (2) the subject scores in the tensor decomposition are a noisy realization of a linear function of \mathbf{Z} . More specifically, we consider the following probabilistic model:

$$x_{itj} = \sum_{k=1}^K u_{ik} \phi_{tk} v_{jk} + \epsilon_{itj}, \quad \epsilon_{itj} \sim \mathcal{N}(0, \sigma_j^2)$$

$$\mathbf{u}_i = (u_{ik})_{k=1}^K \sim \mathcal{N}(\boldsymbol{\mu}_i, \Lambda_f^{-1}), \quad \boldsymbol{\mu}_i = \boldsymbol{\beta}^\top \mathbf{z}_i.$$

Here, (1) $u_{ik}, \phi_{tk}, v_{jk}$ are the subject score, trajectory value and feature loading for factor k in the PARAFAC model and the for observation indexed by (i, t, j) , (2) $\boldsymbol{\beta} \in \mathbb{R}^{q \times K}$ describes the dependence of the expected subject score for subject i on \mathbf{z}_i , and (3) Λ_f is a $K \times K$ diagonal precision matrix. We impose smoothness on the time trajectories $(\phi_{tk})_{t=1}^\top$ and sparsity on $\boldsymbol{\beta}$,

Keywords and phrases: Smooth tensor decomposition, Probabilistic model, Auxiliary covariate.

and estimate these parameters using a mixed EM procedure. More explanation of the model is given in Section 2 and Section 3.

Several related issues arise when estimating the parameters and interpreting the results. We shall summarize them here, with more detailed descriptions in Section 4:

- A good initialization is usually important for iterative optimization of a nonconvex objective. We suggest an initialization procedure combining multilinear SVD (MLSVD) and functional PCA, as an alternative to random parameter initialization that is frequently used in the PARAFAC decomposition.
- Selection of tuning parameters to control the smoothness and sparsity level is non-trivial. It can be computationally expensive to carry out cross-validation. We will suggest a computationally efficient way of auto-tuning in this paper.
- In high dimensions, our estimated model parameters tend to overfit towards the noise, especially in the low signal-to-noise (SNR) setting. We propose a cross-fitting procedure for better reconstructing the signal array, motivated by observations of eigenvalue bias and eigenvector inconsistency in high-dimensional PCA.
- We suggest a computationally efficient method of performing hypothesis testing on β . The constructed p-values can be used as an informative measure of variable importance of Z for patient subgroups.

We provide a python package *spaco* for researchers interested in applying the proposed method.

Our model is built upon a large body of previous literature around PCA, functional PCA, and tensor decomposition. PCA/SVD is the most popular tool for dimension reduction and exploration of multivariate observations. Functional PCA (Besse and Ramsay, 1986; Yao, Müller and Wang, 2005) is often used in particular for modeling longitudinal data. It utilizes the smoothness of time trajectories to handle sparsity in the longitudinal observations, and estimates the eigenvectors and factor scores under this smoothness assumption. The smoothed probabilistic PARAFAC model can be viewed as an extension of functional PCA to the setting of three-way tensor decomposition (Acar and Yener, 2008; Sidiropoulos et al., 2017) using the parallel factor (PARAFAC) model (Harshman and Lundy, 1994; Carroll, Pruzansky and Kruskal, 1980).

Tensor decomposition has been applied to many applications, including speech signal separation (Nion et al., 2009), topic modeling (Anandkumar et al., 2012), communication (Nion and Sidiropoulos, 2010), spectral hypergraph theory and hyper-graph matching (Duchenne et al., 2011), and analysis of biomedical multi-omics data (Taguchi, 2019; Xu et al., 2020). In the study of multivariate longitudinal data in economics (Fan, Fan and Lv, 2008; Lam, Yao and Bathia, 2011; Fan, Liao and Mincheva, 2011), researchers have combined tensor decomposition with auto-cross-covariance estimation and autoregressive models (Wang, Liu and Chen, 2019; Wang et al., 2021). However, these methods usually do not work well with highly sparse data or do not scale well with the feature dimensions, which are important for working with medical data. In Yokota, Zhao and Cichocki (2016) and Imaizumi and Hayashi (2017), the authors introduce smooth tensor decomposition models, and estimate the model parameters by iteratively solving penalized regression problems.

In contrast to the approaches of Yokota, Zhao and Cichocki (2016) and Imaizumi and Hayashi (2017), SPACO is a probabilistic model, which also jointly models the longitudinal data with potentially high-dimensional non-temporal covariates Z . The incorporation of auxiliary covariates can sometimes improve signal estimation quality in the matrix setting (Li, Shen and Huang, 2016), and we observe similar improvements when the auxiliary covariates contain important information about the subject clustering structure. We also suggest methods for good initialization and efficient selection of tuning parameters, which are important in practice but were not discussed in these two previous works.

Even in settings without auxiliary covariates \mathbf{Z} , our combination of probabilistic modeling and automatic parameter tuning provides an attractive alternative to existing smooth tensor decomposition methods. We will refer to this model without covariates as SPACO-, and we will also compare SPACO with SPACO- in our simulations.

The paper is organized as follows. Section 1.1 provides a brief review of the PARAFAC decomposition and related tensor decomposition methods. We describe the SPACO model in Section 2 and model parameter estimation with fixed tuning parameters in Section 3. In Section 4, we discuss parameter initialization, the choice of tuning parameters, cross-fitting, and hypothesis testing. We provide extensive simulations in Section 5 on the different aspects of the SPACO model. Finally, in Section 6, we apply SPACO to two SARS-COV-2 data sets and demonstrate its use in immunological studies.

1.1. Notations and previous work.

1.1.1. *Notations.* Let $\mathbf{X} \in \mathbb{R}^{I \times T \times J}$ be a tensor for some sparse multivariate longitudinal observations, where I is the number of subjects, J is the number of features and T is the number of total unique time points. For any matrix A , we let $A_{i:}/A_{:,i}$ denote its i^{th} row/column, and often write A_i for the i^{th} column $A_{:,i}$. We let $\mathbf{X}_I := (\mathbf{X}_{:,:,1} \cdots \mathbf{X}_{:,:,J}) \in \mathbb{R}^{I \times (TJ)}$, $\mathbf{X}_T := (\mathbf{X}_{:,1}^\top \cdots \mathbf{X}_{:,J}^\top) \in \mathbb{R}^{T \times (IJ)}$, and $\mathbf{X}_J := (\mathbf{X}_{:,1}^\top \cdots \mathbf{X}_{:,T}^\top) \in \mathbb{R}^{J \times (IT)}$ be the matrices created by unfolding \mathbf{X} in the subject/feature/time dimension respectively. We also define:

- Tensor product \odot : for three vectors $a \in \mathbb{R}^I, b \in \mathbb{R}^T, c \in \mathbb{R}^J, A = a \odot b \odot c \in \mathbb{R}^{I \times T \times J}$ with $A_{itj} = a_i b_t c_j$.
- Kronecker product \otimes : for two matrices $A \in \mathbb{R}^{I_1 \times K_1}, B \in \mathbb{R}^{I_2 \times K_2}$,

$$C = A \otimes B = \begin{pmatrix} A_{11}B & \dots & A_{1K_1}B \\ \vdots & \ddots & \vdots \\ A_{I_1 1}B & \dots & A_{I_1 K_1}B \end{pmatrix} \in \mathbb{R}^{(I_1 I_2) \times (K_1 K_2)}.$$

- Column-wise Khatri-Rao product \odot : for two matrices $A \in \mathbb{R}^{I_1 \times K}, B \in \mathbb{R}^{I_2 \times K}, C = A \odot B \in \mathbb{R}^{(I_1 I_2) \times K}$ with $C_{:,k} = (A_{:,k} \otimes B_{:,k})$ for $k = 1, \dots, K$.
- Element-wise multiplication \cdot : for two matrices $A, B \in \mathbb{R}^{I \times K}, C = A \cdot B \in \mathbb{R}^{I \times K}$ with $C_{ik} = (A_{ik} B_{ik})$; for $b \in \mathbb{R}^K, C = A \cdot b = \text{Adiag}\{b_1, \dots, b_K\}$; for $b \in \mathbb{R}^I, C = b \cdot A = \text{diag}\{b_1, \dots, b_I\}A$.
- Inner product $\langle \cdot, \cdot \rangle$: for two vectors $a, b \in \mathbb{R}^I, \langle a, b \rangle = \sum_{i=1}^I a_i b_i$.

1.1.2. *PARAFAC decomposition.* We can decompose a tensor \mathbf{X} based on two factorization models, the PARAFAC/CANDECOMP model (Harshman and Lundy, 1994; Carroll, Pruzansky and Kruskal, 1980) or the Tucker model (Tucker, 1966). The PARAFAC model decomposes \mathbf{X} into a linear combination of rank-one tensors:

$$(1) \quad \mathbf{X} = \sum_{k=1}^K \mathbf{U}_k \odot \Phi_k \odot \mathbf{V}_k.$$

K is the rank of \mathbf{X} and is the minimum number of rank-one tensor needed to produce \mathbf{X} . Unlike SVD, the above tensor decomposition does not require the orthogonality of $\mathbf{U}/\Phi/\mathbf{V}$.

Low-rank approximation of an observed tensor \mathbf{X} is often performed by solving the optimization problem

$$(2) \quad \min_{\mathbf{U}, \Phi, \mathbf{V}} \|\mathbf{X} - \sum_{k=1}^K [\mathbf{U}_k \odot \Phi_k \odot \mathbf{V}_k]\|_F^2,$$

for small K . This low-rank approximation is NP-hard to compute in general (Johan, 1990), and a common heuristic is to find a locally optimal solution by updating U, Φ, V with alternating least squares (ALS):

$$\begin{aligned} U &\leftarrow \arg \min_U \|X_I - U(V \odot \Phi)^\top\|_F^2 \\ \Phi &\leftarrow \arg \min_\Phi \|X_T - \Phi(V \odot U)^\top\|_F^2 \\ V &\leftarrow \arg \min_V \|X_J - V(\Phi \odot U)^\top\|_F^2 \end{aligned}$$

Another tensor decomposition model is the Tucker model. This usually requires (U, Φ, V) to be orthogonal matrices, and introduces a cross-component correlation using a non-diagonal core array $G \in \mathbb{R}^{K \times K \times K}$. Solutions from the PARAFAC and Tucker decompositions can be expressed using one another (ignoring computational challenges). We omit details of the Tucker model and refer interested readers to Sidiropoulos et al. (2017).

1.1.3. *Tensor decomposition for multivariate longitudinal data.* Tensor decomposition has been proposed previously to analyze multivariate longitudinal data. In Wang et al. (2021), the authors consider an autoregressive time-series model where

$$x_t = A_1 x_{t-1} + \dots + A_p x_{t-p} + \epsilon_t,$$

with $x_t = (x_{itj} : i = 1, \dots, I, j = 1, \dots, J)$ and the coefficient matrix (A_1, \dots, A_p) assumed to be an unfolding of a tensor with multi-linear low rank (Tucker, 1966). In Wang, Liu and Chen (2019), the authors consider a procedure based on auto-cross-covariance estimation. Setting $x_{t,j} = (x_{itj} : i = 1, \dots, I)$, the authors first estimate the empirical auto-cross-covariance

$$\hat{\Omega}_{x,j_1,j_2}(h) = \frac{1}{T-h} \sum_{t=1}^{T-h} x_{t,j_1} x_{t+h,j_2}^\top,$$

followed by a low-rank decomposition of $M = \sum_{h=1}^{h_0} \sum_{j_1=1}^p \sum_{j_2=1}^p \hat{\Omega}_{x,j_1,j_2}(h)$ to estimate the auto-covariance of the structure of latent time trajectories. Both these two approaches do not deal well with sparsity along the time dimension. The second approach also scales poorly with the feature dimension J , which can be hundreds or thousands in medical applications.

Tensor decomposition models that impose smoothness along the time dimension were studied in Imaizumi and Hayashi (2017) and Yokota, Zhao and Cichocki (2016). In Imaizumi and Hayashi (2017), the authors introduced smoothness to the reconstructed array \hat{X} , whereas in Yokota, Zhao and Cichocki (2016), the authors considered a PARAFAC model with a smoothness penalty on the decomposed components. Both works proposed optimization approaches for parameter estimation.

We extend the approach of Yokota, Zhao and Cichocki (2016) by considering a probabilistic model, and incorporate the auxiliary covariates Z . A probabilistic model can be helpful with highly sparse and irregular longitudinal sampling, by modeling randomness in the observations. Auxiliary covariates are common in medical applications where biologists and clinicians may want to understand whether age, sex, disease status, or other personalized medical histories are related to subjects' temporal trajectories.

2. SPACO model.

2.1. *A probabilistic PARAFAC decomposition with covariates.* Our probabilistic PARAFAC model extends probabilistic matrix factorization (Tipping and Bishop, 1999; Mnih and Salakhutdinov, 2007; Hinrich and Mørup, 2019) to higher-order arrays. In the probabilistic PARAFAC, we assume \mathbf{X} to be a noisy realization of an underlying signal array $\mathbf{F} = \sum_{k=1}^K \mathbf{U}_k \odot \mathbf{\Phi}_k \odot \mathbf{V}_k$. We stack $\mathbf{U}_k/\mathbf{\Phi}_k/\mathbf{V}_k$ as the columns of $\mathbf{U}/\mathbf{\Phi}/\mathbf{V}$, denote the rows of $\mathbf{U}/\mathbf{\Phi}/\mathbf{V}$ by $\mathbf{u}_i/\phi_t/v_j$, and their entries by $u_{ik}/\phi_{tk}/v_{jk}$. We let x_{itj} denote the (i, t, j) -entry of \mathbf{X} .

Even though \mathbf{X} is often of high rank, we consider the scenario where the rank of \mathbf{F} is small. The mathematical formulation of the probabilistic model is given by

$$(3) \quad \begin{aligned} x_{itj} &= \sum_{k=1}^K u_{ik} \phi_{tk} v_{jk} + \epsilon_{itj}, \\ \mathbf{u}_i &\sim \mathcal{N}(\boldsymbol{\mu}_i, \Lambda_f^{-1}), \\ \epsilon_{ijt} &\sim \mathcal{N}(0, \sigma_j^2), \end{aligned}$$

where $\Lambda_f = \text{diag}\{\frac{1}{s_1^2}, \dots, \frac{1}{s_K^2}\}$ is a $K \times K$ diagonal precision matrix. Without covariates, we set the mean parameter $\boldsymbol{\mu}_i = \mathbf{0}$. If we are interested in explaining the heterogeneity in $\boldsymbol{\mu}_i$ across subjects with auxiliary covariates $\mathbf{Z} \in \mathbb{R}^{I \times q}$, then we may $\boldsymbol{\mu}_i$ as a function of $\mathbf{z}_i := \mathbf{Z}_{i,:}$. Here, we consider a linear model

$$\mu_{ik} = \mathbf{z}_i^\top \boldsymbol{\beta}_k, \quad \forall k = 1, \dots, K.$$

To avoid confusion, we will always call \mathbf{X} the ‘‘features’’, and \mathbf{Z} the ‘‘covariates’’ or ‘‘variables’’.

The component \mathbf{U} characterizes differences across subjects and can be used for subject clustering. We refer to \mathbf{U} as the subject scores. The component \mathbf{V} reveals the composition of the factors in terms of the original features, and can be used for downstream interpretation. We refer to \mathbf{V} as the feature loadings. The component $\mathbf{\Phi}$ can be interpreted as function values sampled from some underlying smooth functions $\phi_k(t)$ at a set of discrete-time points, so that

$$\mathbf{\Phi}_k = (\phi_k(t_1), \dots, \phi_k(t_T)).$$

We refer to $\mathbf{\Phi}$ as the time trajectories. The smoothness assumption enables us to work with sparse longitudinal observations efficiently. We treat \mathbf{U} as latent variables and let

$$\Theta = \{\mathbf{V}, \mathbf{\Phi}, \boldsymbol{\beta}, (\sigma_j^2, j = 1, \dots, J), (s_k^2, k = 1, \dots, K)\}$$

denote the set of model parameters.

If \mathbf{U} is completely observed, then it is straightforward to write down the complete-data log-likelihood: Recalling that $\mathbf{X}_I \in \mathbb{R}^{I \times (TJ)}$ is the unfolding of \mathbf{X} in the subject direction, we write \vec{i} for the indices of observed values in the i^{th} row of \mathbf{X}_I , and $\mathbf{X}_{I, \vec{i}}$ for the vector of these observed values. Each such observed value x_{itj} has noise variance σ_j^2 , and we write $\Lambda_{\vec{i}}$ to represent the diagonal precision matrix of values $1/\sigma_j^2$ corresponding to the indices in \vec{i} . Similarly, we define $\{\vec{t}, \mathbf{X}_{T, \vec{t}}, \Lambda_{\vec{t}}\}$ for the unfolding $\mathbf{X}_T \in \mathbb{R}^{T \times (IJ)}$, and $\{\vec{j}, \mathbf{X}_{J, \vec{j}}, \Lambda_{\vec{j}}\}$ for $\mathbf{X}_J \in \mathbb{R}^{J \times (IT)}$.

The complete data log-likelihood is given by

$$(4) \quad \begin{aligned} L(\mathbf{X}, \mathbf{U} | \Theta) &= -\frac{1}{2} \sum_i \left(\mathbf{X}_{I, \vec{i}} - (\mathbf{V} \odot \mathbf{\Phi})_{\vec{i}} \mathbf{u}_i \right)^\top \Lambda_{\vec{i}} \left(\mathbf{X}_{I, \vec{i}} - (\mathbf{V} \odot \mathbf{\Phi})_{\vec{i}} \mathbf{u}_i \right)^\top \\ &\quad - \frac{1}{2} \sum_i (\mathbf{u}_i - \boldsymbol{\beta}^\top \mathbf{z}_i)^\top \Lambda_f (\mathbf{u}_i - \boldsymbol{\beta}^\top \mathbf{z}_i) + \frac{1}{2} \sum_i (\log |\Lambda_{\vec{i}}| + I \log |\Lambda_f|). \end{aligned}$$

In practice, \mathbf{U} is not observed. As a result, we consider the marginal log-likelihood integrating out \mathbf{U} in our model. Under the Gaussian assumption for \mathbf{U} , this marginalization results in the closed-form expression

$$(5) \quad L(\mathbf{X}|\Theta) \propto -\frac{1}{2} \sum_i \left(\mathbf{X}_{I,i}^\top \Lambda_{\vec{i}} \mathbf{X}_{I,i} + \mathbf{z}_i^\top \beta \Lambda_f \beta^\top \mathbf{z}_i - \boldsymbol{\mu}_i^\top \Sigma_i^{-1} \boldsymbol{\mu}_i \right) + \frac{1}{2} \left(\sum_i (\log |\Lambda_{\vec{i}}| + \log |\Sigma_i|) + I \log |\Lambda_f| \right),$$

where

$$(6) \quad \Sigma_i = \left(\Lambda_f + (\mathbf{V} \odot \Phi)_i^\top \Lambda_{\vec{i}} (\mathbf{V} \odot \Phi)_i \right)^{-1}, \quad \boldsymbol{\mu}_i = \Sigma_i \left(\Lambda_f \beta^\top \mathbf{z}_i + (\mathbf{V} \odot \Phi)_i^\top \Lambda_{\vec{i}} \mathbf{X}_{I,i} \right)$$

are the posterior covariance and mean of \mathbf{u}_i . We discuss identifiability of model parameters and smoothness and sparsity penalties in the remainder of this section. We give a mixed EM-procedure for greedily maximizing the penalized likelihood in Section 3.

2.2. Identifiability of model parameters. The parameters of the model described in Section 2.1 are not strictly identifiable, for the following reasons:

- Permuting the factors $\{1, \dots, K\}$ yields the same model. We will always order the factors such that $s_1^2 \geq \dots \geq s_K^2$ when reporting results.
- The rescaling $(c_1 \mathbf{U}_k, c_2 \Phi_k, c_3 \mathbf{V}_k)$ represents the same model as $(\mathbf{U}_k, \Phi_k, \mathbf{V}_k)$ if $c_1 c_2 c_3 = 1$. Thus, in our setting where \mathbf{U} are latent variables, the parameters $(\Phi_k, \mathbf{V}_k, \beta_k, s_k^2)$ represent the same model as $(c_1 \Phi_k, c_2 \mathbf{V}_k, c_3 \beta_k, c_3^2 s_k^2)$:

PROPOSITION 2.1. *Let $\Theta_k(c_1, c_2, c_3) = \{c_1 \Phi_k, c_2 \mathbf{V}_k, c_3 \beta_k, c_3^2 s_k^2\}$, and let Θ_{k^c} be the remaining parameters of Θ . Then*

$$L(\mathbf{X}|\Theta_k(c_1, c_2, c_3), \Theta_{k^c}) = L(\mathbf{X}|\Theta_k(1, 1, 1), \Theta_{k^c}), \text{ for all } c_1 c_2 c_3 = 1.$$

This non-identifiability may cause problems in algorithmic convergence, and the specific scaling of Φ_k is also important for the smoothness penalization that we later introduce. To avoid these problems, we will enforce the scalings

$$(7) \quad \|\mathbf{V}_k\|_2^2 = 1, \quad \|\Phi_k\|_2^2 = T.$$

2.3. Smoothness and sparsity. We assume that the time component Φ_k is sampled from a slowly varying trajectory function $\phi_k(t)$. We encourage smoothness of $\phi_k(t)$ by directly penalizing the function values via a penalty term $\sum_k \lambda_{1k} \Phi_k^\top \Omega \Phi_k$.

There are different types of penalty matrices Ω that practitioners can use. In this paper, we consider a Laplacian smoothing (Sorkine et al., 2004) with weighted adjacency matrix Γ :

$$\Omega = \Gamma^\top \Gamma, \quad \Gamma = \begin{pmatrix} \frac{1}{\mathcal{T}[2]-\mathcal{T}[1]} - \frac{1}{\mathcal{T}[2]-\mathcal{T}[1]} & \cdots & 0 & & 0 \\ 0 & \frac{1}{\mathcal{T}[3]-\mathcal{T}[2]} & \cdots & 0 & \vdots \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & \frac{1}{\mathcal{T}[T]-\mathcal{T}[T-1]} - \frac{1}{\mathcal{T}[T]-\mathcal{T}[T-1]} & \end{pmatrix} \in \mathbb{R}^{T \times (T-1)}$$

If practitioners want $\phi_k(t)$ to have slowly varying derivatives, they can also use a penalty matrix that penalizes changes in gradients over time.

When the number of covariates q in \mathbf{Z} is moderately large, we may wish to impose sparsity in the β parameter. We encourage such sparsity by including a lasso penalty (Tibshirani,

2011) in the model. Our goal is then to find parameters maximizing the expected penalized log-likelihood, or minimizing the penalized expected deviance loss, under the scaling constraints of eq. (7):

$$\begin{aligned} \min J(\Theta) &:= -\frac{1}{2}L(\mathbf{X}|\Theta) + \sum_{k=1}^K \lambda_{1k} \mathbf{\Phi}_k^\top \Omega \mathbf{\Phi}_k + \sum_k \lambda_{2k} |\beta_k| \\ (8) \quad \text{s.t.} \quad &\|\mathbf{V}_k\|_2^2 = 1, \|\mathbf{\Phi}_k\|_2^2 = T, \text{ for all } k = 1, \dots, K. \end{aligned}$$

Eq. (8) describes a non-convex problem. However, we can find locally optimal solutions via an alternating update procedure. We may update β_k using lasso regression, fixing other parameters. However, there is no simple subroutine for updating parameters like \mathbf{V}_k , even with other parameters fixed. This leads to a mixed EM procedure where we update β_k maximizing the marginal log-likelihood, and update other parameters using the EM algorithm. We give details of our iterative estimation procedure in Section 3 and show that the loss is nonincreasing over the iterations.

3. Parameter estimation. Given penalty terms λ_{1k} , λ_{2k} , we alternately update parameters β , \mathbf{V} , \mathbf{U} , s^2 and σ^2 with a mixed EM procedure described in Algorithm 1. We briefly explain the updating steps here:

- Given other parameters, we can minimize over β by solving a least-squares regression problem with lasso penalty.
- Fixing β , we update the other parameters using an EM procedure. Denote the current parameters as Θ_0 . Our goal is to maximize the penalized expected log-likelihood

$$\mathbb{E}_{\mathbf{U}|\Theta_0} \left\{ L(\mathbf{X}, \mathbf{U}|\Theta) - \sum_k \lambda_{1k} \mathbf{\Phi}_k^\top \Omega \mathbf{\Phi}_k - \sum_k \lambda_{2k} |\beta_k| \right\},$$

under the posterior distribution $\mathbf{U}|\Theta_0$. However, this is still a non-convex problem under the constraints $\|\mathbf{V}_k\|_2^2 = 1$ and $\|\mathbf{\Phi}_k\|_2^2 = T$.

- In Algorithm 1, steps 6-15, we update parameters \mathbf{V}_k , $\mathbf{\Phi}_k$, s_k^2 and σ_j^2 iteratively. In particular, due to the nonconvex constraints on \mathbf{V}_k and $\mathbf{\Phi}_k$, we update these via proximal gradient descent. We define proximal mappings $\text{prox}_{\mathbf{V}_k}(\mathbf{y}; \mathbf{y}_0, \rho) / \text{prox}_{\mathbf{\Phi}_k}(\mathbf{y}; \mathbf{y}_0, \rho)$ for updating $\mathbf{V}_k / \mathbf{\Phi}_k$ as

$$\begin{aligned} \text{prox}_{\mathbf{V}_k}(\mathbf{y}; \mathbf{y}_0, \rho) &= \begin{cases} \frac{1}{2} \|\mathbf{y} - \mathbf{y}_0\|_2^2 - \rho \mathbf{y}^\top \frac{\partial \mathbb{E}_{\mathbf{U}|\Theta_0} L(\mathbf{X}, \mathbf{U}|\Theta)}{\partial \mathbf{V}_k} \Big|_{\mathbf{V}_k = \mathbf{y}_0} & \text{if } \|\mathbf{y}\|_2^2 = 1 \\ \infty & \text{otherwise,} \end{cases} \\ \text{prox}_{\mathbf{\Phi}_k}(\mathbf{y}; \mathbf{y}_0, \rho) &= \begin{cases} \frac{1}{2} \|\mathbf{y} - \mathbf{y}_0\|_2^2 - \rho \mathbf{y}^\top \left(\frac{\partial \mathbb{E}_{\mathbf{U}|\Theta_0} L(\mathbf{X}, \mathbf{U}|\Theta)}{\partial \mathbf{\Phi}_k} \Big|_{\mathbf{\Phi}_k = \mathbf{y}_0} - \lambda_{1k} \Omega \mathbf{y}_0 \right) & \text{if } \|\mathbf{y}\|_2^2 = T \\ \infty & \text{otherwise.} \end{cases} \end{aligned}$$

In steps 7-9 and 11-13, we perform a proximal update for \mathbf{V}_k and $\mathbf{\Phi}_k$ until no improvement is observed in our objective.

Algorithm 1: SPACO with fixed penalties

Data: $\mathbf{X}, \Omega, \lambda_1, \lambda_2, K$
Result: Estimated $\mathbf{V}, \Phi, \beta, s^2, \sigma^2$ and $P(\mathbf{U}|\Theta, \mathbf{X})$, the posterior distribution of \mathbf{U} .

- 1 Initialization of $\mathbf{V}, \Phi, \beta, s^2, \sigma^2$;
- 2 **while** *Not converged* **do**
- 3 **for** $k = 1, \dots, K$ **do**
- 4 $\beta_{:,k} \leftarrow \arg \max_{\beta_{:,k}} \{L(\mathbf{X}|\Theta) + \lambda_{2k} |\beta_{:,k}|\}$
- 5 **end**
- 6 For $k = 1, \dots, K$:
- 7 **while** *Not converged* **do**
- 8 $\mathbf{V}_k \leftarrow \arg \min_{\mathbf{y}} \text{prox}_{\mathbf{V}_k}(\mathbf{y}; \mathbf{V}_k, \rho)$ for a sufficiently small ρ .
- 9 **end**
- 10 For $k = 1, \dots, K$:
- 11 **while** *Not converged* **do**
- 12 $\Phi_k \leftarrow \arg \min_{\mathbf{y}} \text{prox}_{\Phi_k}(\mathbf{y}; \Phi_k, \rho)$ for a sufficiently small ρ .
- 13 **end**
- 14 For $k = 1, \dots, K$: $s_k^2 \leftarrow \arg \min_{s_k^2} \mathbb{E}_{\mathbf{U}|\Theta_0} L(\mathbf{X}, \mathbf{U}|\Theta)$.
- 15 For $j = 1, \dots, J$: $\sigma_j^2 \leftarrow \arg \min_{\sigma_j^2} \mathbb{E}_{\mathbf{U}|\Theta_0} L(\mathbf{X}, \mathbf{U}|\Theta)$.
- 16 **end**

In more detail, the following are the explicit forms of the updates of all model parameters: Recall the posterior covariance Σ_i and mean μ_i from eq. (6) given the current model parameters. In the context of two indices j and t , recall also the index vectors \vec{j} and \vec{t} of observed columns in the j^{th} and t^{th} rows of \mathbf{X}_J and \mathbf{X}_T . Stack μ_1, \dots, μ_I as the rows of $\mu \in \mathbb{R}^{I \times K}$, and let $(\Phi \odot \mu)_{\vec{j}}$ and $(\mathbf{V} \odot \mu)_{\vec{t}}$ denote these corresponding rows of $(\Phi \odot \mu)$ and $(\mathbf{V} \odot \mu)$. Finally, let $O_{it} = 1$ if features $\mathbf{X}_{i,t,:}$ are observed, and $O_{it} = 0$ otherwise. The exact parameter update rules are provided in Lemma 3.1.

LEMMA 3.1. *The parameter update steps in Algorithm 1 take the following forms:*

- In line 4, $\beta_k = \arg \max_{\beta_k} \{L(\mathbf{X}|\Theta) + \lambda_{2k} |\beta_k|\} = \arg \min_{\beta_k} \frac{1}{2} \|\tilde{\mathbf{y}} - \tilde{\mathbf{Z}} \beta_k\|_2^2 + \lambda_{2k} |\beta_k|$, where

$$(9) \quad \tilde{\mathbf{z}}_i = \sqrt{\frac{1}{s_k^2} \frac{s_k^2 - (\Sigma_i)_{k,k}}{s_k^2}} \mathbf{z}_i,$$

$$(10) \quad \tilde{\mathbf{y}}_i = \sqrt{\frac{1}{s_k^2} \frac{s_k^2 - (\Sigma_i)_{k,k}}{s_k^2 - (\Sigma_i)_{k,k}}} \left((\Sigma_i)_{k,:} (\mathbf{V} \odot \Phi)_{\vec{t}}^\top \Lambda_{\vec{t}}^\top \mathbf{X}_{I,\vec{i}} + (\Sigma_i)_{k,k^c} \Lambda_{f,k^c} \beta_{k^c}^\top \mathbf{z}_i \right).$$

Here $(\Sigma_i)_{k,:}$ is the k^{th} row of Σ_i , and $(\Sigma_i)_{k,k^c}$ is this row with k^{th} entry removed; Λ_{f,k^c} is the sub matrix of Λ_f with the k^{th} column and row removed.

- Define $M^j = \frac{1}{\sigma_j^2} \left[(\Phi \odot \mu)_{\vec{j}} (\Phi \odot \mu)_{\vec{j}}^\top + (\sum_t (\sum_{i:O_{it}=1} \Sigma_i) \cdot (\phi_t \phi_t^\top)) \right] \in \mathbb{R}^{K \times K}$, $h_j = \frac{1}{\sigma_j^2} (\Phi \odot \mu)_{\vec{j},k}^\top \mathbf{X}_{J,\vec{j}}$. Let $A_k^V = \text{diag}\{M_{k,k}^1, \dots, M_{k,k}^J\}$, $a \in \mathbb{R}^J$ with $a_j = h_j - \mathbf{V}_{k^c} M_{k^c,k}^j$, where \mathbf{V}_{k^c} is \mathbf{V} excluding its k^{th} column. Then line 8 is

$$\mathbf{V}_k = \frac{\tilde{\mathbf{V}}_k}{\|\tilde{\mathbf{V}}_k\|_2}, \quad \tilde{\mathbf{V}}_k = \mathbf{V}_k + \rho(a - A_k^V \mathbf{V}_k).$$

- Define $M^t = \left[(\mathbf{V} \odot \mu)_{\vec{t}}^\top \Lambda_{\vec{t}} (\mathbf{V} \odot \mu)_{\vec{t}} + \sum_j \frac{1}{\sigma_j^2} (\sum_{i:O_{it}=1} \Sigma_i) \cdot (\mathbf{v}_j \mathbf{v}_j^\top) \right]$, $h_t = (\mathbf{V} \odot \mu)_{\vec{t},k}^\top \Lambda_{\vec{t}}^\top \mathbf{X}_{T,\vec{t}}$. Let $A_k^\Phi = \text{diag}\{M_{k,k}^1, \dots, M_{k,k}^T\}$, $a \in \mathbb{R}^T$ with $a_t = h_t - \Phi_{k^c} M_{k^c,k}^t$. Then

line 12 is

$$\mathbf{\Phi}_k = \frac{\tilde{\mathbf{\Phi}}_k}{\|\tilde{\mathbf{\Phi}}_k\|_2}, \quad \tilde{\mathbf{\Phi}}_k = \mathbf{\Phi}_k + \rho(a - A_k^\mathbf{\Phi} \mathbf{\Phi}_k - \lambda_{1k} \mathbf{\Phi}_k).$$

• In lines 14 and 15,

$$s_k^2 \leftarrow \frac{1}{I} \sum_i \left[(\mu_{ik} - \mathbf{z}_i^\top \boldsymbol{\beta}_k)^2 + (\Sigma_i)_{kk} \right]$$

and

$$\sigma_j^2 \leftarrow \frac{1}{|\vec{j}|} \left[(X_{J,\vec{j}} - (\mathbf{\Phi} \odot \boldsymbol{\mu})_{\vec{j}} \mathbf{v}_j)^\top (X_{J,\vec{j}} - (\mathbf{\Phi} \odot \boldsymbol{\mu})_{\vec{j}} \mathbf{v}_j) + \mathbf{v}_j^\top \left(\sum_t \left(\sum_{i:O_{it}=1} \Sigma_i \right) \cdot (\phi_t \phi_t^\top) \right) \mathbf{v}_j \right],$$

where $|\vec{j}|$ is the total number of observed entries for feature j across all subjects and times.

Even though $\|\mathbf{V}_k\|_2^2 = 1$ and $\|\mathbf{\Phi}_k\|_2^2 = T$ are non-convex constraints, Algorithm 1 guarantees that the loss is non-increasing as long as proximal gradient descent uses a sufficiently small step size when updating \mathbf{V}_k and $\mathbf{\Phi}_k$. Let $\text{eig}_{\max}(A)$ be the largest eigenvalue of a matrix A .

THEOREM 3.2. *Let $\rho_k(\mathbf{V})$ and $\rho_k(\mathbf{\Phi})$ be the step sizes in proximal gradient descent for \mathbf{V}_k and $\mathbf{\Phi}_k$. Let Θ^ℓ be parameters at the beginning of the ℓ^{th} iteration of the outer **while** loop and $\Theta^{\ell+1}$ be the parameter values at the end of this iteration. When $\frac{1}{\rho_k(\mathbf{V})} \geq \max_j M_{kk}^j$, $\frac{1}{\rho_k(\mathbf{\Phi})} \geq \text{eig}_{\max}(A_k^\mathbf{\Phi} + \lambda_{1k} \Omega)$, we have $J(\Theta^\ell) \geq J(\Theta^{\ell+1})$.*

Table 1 tracks the computational cost at each of the updating steps. (Here $|\vec{i}|$ is the number of observed entries in \vec{i} , etc.)

TABLE 1
Computational complexity for SPACO

Update posteriors for \mathbf{U}	$\mathcal{O}(K^2 \vec{i})$ for Σ_i and μ_i , and $\mathcal{O}(K^2 J \sum_{it} O_{it})$ in total.
Update $\mathbf{\Phi}$	$\mathcal{O}(K^2 \vec{t})$ for M^t, h_t , $\mathcal{O}(T^2 m)$ for m proximal updates, $\mathcal{O}(K^2 J \sum_{it} O_{it} + KT^2 m)$ in total.
Update \mathbf{V}	$\mathcal{O}(K^2 \vec{j})$ for M^j, h_j , $\mathcal{O}(Jm)$ for m proximal updates, $\mathcal{O}(K^2 J \sum_{it} O_{it} + KJm^2)$ in total.
Update $\boldsymbol{\beta}$	$\mathcal{O}(mKIqq^*)$ where $q^* = \text{sparsity level}$, $m = \text{number of iterations for lasso to converge}$.
Update σ^2	$\mathcal{O}(KJ \sum_{it} O_{it})$
Update s^2	$\mathcal{O}(IqK)$

Treating the number of lasso and proximal descent iterations as fixed, the computational complexity for each full update is $\mathcal{O}(KT^2 + K^2 J \sum_{it} O_{it} + K I q q^*)$, which is at most linear in the total number of observations $J \sum_{it} O_{it}$ and quadratic in the number of unique time points. Hence, SPACO scales well with the problem dimensions.

4. Initialization, tuning, cross-fitting, and variable importance.

4.1. *Model parameter initialization.* Random initialization for the PARAFAC model is often used, for example in the smooth PARAFAC decomposition of [Yokota, Zhao and Cichocki \(2016\)](#). An alternative approach is to first perform Tucker decomposition using HOSVD/MLSVD ([De Lathauwer, De Moor and Vandewalle, 2000](#)), followed by PARAFAC decomposition on the small core tensor $\mathbf{G} \in \mathbb{R}^{K \times K \times K}$ ([Bro and Andersson, 1998](#); [Phan,](#)

Tichavský and Cichocki, 2013). We propose an efficient initialization for SPACO based on this approach, as described in Algorithm 2. Our approach follows a procedure similar to (Bro and Andersson, 1998; Phan, Tichavský and Cichocki, 2013), where we first form a Tucker decomposition using a MLSVD-type procedure $\mathbf{X} \approx [\mathbf{U}_\perp, \mathbf{\Phi}_\perp, \mathbf{V}_\perp; \mathbf{G}]$ with orthonormal $\mathbf{U}_\perp \in \mathbb{R}^{I \times K}$, $\mathbf{\Phi}_\perp \in \mathbb{R}^{T \times K}$ and $\mathbf{V}_\perp \in \mathbb{R}^{J \times K}$. Our initialization is different from MLSVD because we account for missing data and smoothness in $\mathbf{\Phi}$. Algorithm 2 consists of the following steps:

- Same as in MLSVD, we perform SVD on \mathbf{X}_J to get \mathbf{V}_\perp .
- We project \mathbf{X}_J onto each column of \mathbf{V}_\perp and perform functional PCA on a total product matrix, to estimate $\mathbf{\Phi}_\perp$.
- We run a ridge-penalized regression of rows of \mathbf{X}_I on $\mathbf{V}_\perp \otimes \mathbf{\Phi}_\perp$, and estimate \mathbf{U}_\perp and \mathbf{G} from the regression coefficients.

Algorithm 2: Initialization of SPACO

- 2 Perform SVD on \mathbf{X}_J using observed columns, and let \mathbf{V}_\perp be the top K left singular vectors.
- 4 Set $\mathbf{Y}(k) = (\mathbf{Y}_1(k), \dots, \mathbf{Y}_T(k)) \in \mathbb{R}^{I \times T}$, where $\mathbf{Y}_t(k) = \mathbf{X}_{:,t,:}(\mathbf{V}_\perp)_k \in \mathbb{R}^I$ is the projection along the k^{th} column of \mathbf{V}_\perp of the data at time t .
- 6 Perform functional PCA by smooth estimation of the total product matrix of $\mathbf{Y}(k)$:

$$\mathbf{W} = \sum_{k=1}^K \mathbf{W}(k) \in \mathbb{R}^{T \times T}, \quad \mathbf{W}_{t,t'}(k) = \mathbb{E}[\mathbf{Y}_t(k)\mathbf{Y}_{t'}(k)].$$

Here $\mathbb{E}[\mathbf{Y}_t(k)\mathbf{Y}_{t'}(k)]$ is the same product across subjects $i = 1, \dots, I$. Smooth estimation of \mathbf{W} is done via a local linear regression estimate described in Yao, Müller and Wang (2005). We provide details in Appendix B.1 for the sake of completeness. Let $\mathbf{\Phi}_\perp$ be the top K eigenvectors of the smooth estimate of \mathbf{W} .

- 8 Let $\tilde{\mathbf{U}} = \arg \min_{\mathbf{U}} \{ \|\mathbf{X}_I - \mathbf{U}(\mathbf{V}_\perp \otimes \mathbf{\Phi}_\perp)^\top\|_F^2 + \delta \|\mathbf{U}\|_F^2 \} \in \mathbb{R}^{I \times K^2}$, where δ is a small regularization parameter to avoid severe over-fitting to the noise. By default, we set $\delta = \frac{1}{\sqrt{J \times T}}$.
- 10 Let \mathbf{U}_\perp be the top K left singular eigenvectors of $\tilde{\mathbf{U}}$, and $\tilde{\mathbf{G}} = \mathbf{U}_\perp^\top \tilde{\mathbf{U}} \in \mathbb{R}^{K \times K^2}$. Let $\mathbf{G} \in \mathbb{R}^{K \times K \times K}$ be the estimated core array from rearranging $\tilde{\mathbf{G}}$ with

$$\mathbf{G}_{k,:,:} = \begin{pmatrix} \tilde{\mathbf{G}}_{k,1:K}^\top \\ \vdots \\ \tilde{\mathbf{G}}_{k,K:(K-1):K^2}^\top \end{pmatrix}.$$

- 12 Let $\sum_{k=1}^K \mathbf{A}_k \otimes \mathbf{B}_k \otimes \mathbf{C}_k$ be the rank-K PARAFAC approximation of \mathbf{G} . Stack these as the columns of $\mathbf{A}, \mathbf{B}, \mathbf{C} \in \mathbb{R}^{K \times K}$, and set

$$[\mathbf{U}, \mathbf{\Phi}, \mathbf{V}] = [\mathbf{U}_\perp \mathbf{A}, \mathbf{\Phi}_\perp \mathbf{B}, \mathbf{V}_\perp \mathbf{C}].$$

- 13 For each $k = 1, \dots, K$, rescale the columns $(\mathbf{U}_k, \mathbf{\Phi}_k, \mathbf{V}_k) \mapsto (c_1 \mathbf{U}_k, c_2 \mathbf{V}_k, c_3 \mathbf{V}_k)$ to satisfy the constraints of eq. (7), where $c_1 c_2 c_3 = 1$. These are our initializations for $\mathbf{V}, \mathbf{\Phi}$, and we initialize $\boldsymbol{\beta} = \mathbf{0}$ and s_k^2 as the sample variance of \mathbf{U}_k .
-

In a noiseless model with $\epsilon = 0$ and complete temporal observations, one may replace the functional PCA step of Algorithm 2 with standard PCA. Then $[\mathbf{U}, \mathbf{\Phi}, \mathbf{V}]$ becomes a PARAFAC decomposition of $\frac{1}{1+\delta} \mathbf{X}$.

LEMMA 4.1. *Suppose $\mathbf{X} = \sum_{k=1}^K \mathbf{U}_k^* \odot \Phi_k^* \odot \mathbf{V}_k^*$ and is completely observed. Replace Φ_{\perp} in Algorithm 2 by the top K eigenvectors of $\mathbf{W} = \frac{1}{T} \sum_{k=1}^K \mathbf{Y}(k)^{\top} \mathbf{Y}(k)$. Then, the output $\mathbf{U}, \Phi, \mathbf{V}$ of Algorithm 2 forms a PARAFAC decomposition of $\frac{1}{1+\delta} \mathbf{X}$: $\mathbf{X} = (1 + \delta) \sum_{k=1}^K \mathbf{U}_k \odot \Phi_k \odot \mathbf{V}_k$.*

In Algorithm 2, we estimate Φ_{\perp} based on functional PCA to account for the temporal sparsity and observational noise. When the data are incomplete, setting $\delta > 0$ can also help avoid estimates driven by subjects exhibiting low signal to noise ratio, i.e., subjects i where $(\mathbf{V}_{\perp} \otimes \Phi_{\perp})_{i,:}^{\top} (\mathbf{V}_{\perp} \otimes \Phi_{\perp})_{i,:} \in \mathbb{R}^{K^2 \times K^2}$ is rank-deficient or poorly conditioned.

4.2. *Tuning Parameters.* One way to choose the tuning parameters λ_{1k} and λ_{2k} is to use cross-validation. However, this can be computationally expensive even if we tune each parameter sequentially. It is also difficult to determine a good set of candidate values for the parameters before running SPACO. We propose instead to determine the tuning parameters via nested cross-validation, which has been shown to be empirically useful (Huang et al., 2008; Li, Shen and Huang, 2016). In nested cross-validation, the parameters are tuned within their corresponding subroutines. Although this is not as disciplined as full cross-validation, it incurs much less computational cost.

More specifically, we propose the following procedure:

- In the update for Φ_k , perform column-wise leave-one-out cross-validation to select λ_{1k} , where we leave out all observations from a single time point.
- In the update for $\beta_{:,k}$, perform K -fold cross-validation to select λ_{2k} .

We provide details about the column-wise leave-one-out cross-validation in Appendix B.2.

4.3. *Cross-fitting.* We now discuss reconstruction of the underlying signal tensor $\mathbf{F} = \sum_{k=1}^K \mathbf{U}_k \odot \Phi_k \odot \mathbf{V}_k$. Under the probabilistic model of SPACO, we can perform this reconstruction using the estimated posterior means μ and Φ, \mathbf{V} :

$$\hat{\mathbf{F}} = \sum_{k=1}^K \mu_k \odot \Phi_k \odot \mathbf{V}_k.$$

However, in the two-dimensional matrix setting, it is known that such an estimate may be inaccurate due to eigenvalue bias and eigenvector inconsistency in high dimensions (Baik and Silverstein, 2006; Paul, 2007; Mestre, 2008), and shrinkage of the estimated factor scores can often reduce the MSE. The same phenomena exist in the three-dimensional tensor setting of SPACO. Here, we propose a shrinkage on μ based on cross-fitting.

Let $\mu \in \mathbb{R}^{I \times K}$ be the posterior means (from either SPACO with covariates or SPACO-without). We perform the reconstruction by shrinking each μ_k by a constant $c_k \in [0, 1]$, and use $\tilde{\mu}_k = c_k \mu_k$ in the reconstruction. Our determination of c_k is motivated by the following toy example.

EXAMPLE. Consider a matrix that is the sum of a rank-one component and mean-zero observational noise, $\mathbf{X} = s\mathbf{u}\mathbf{v}^{\top} + \epsilon$ where $s > 0$ and \mathbf{u}, \mathbf{v} are unit-norm. Let $\hat{s}, \hat{\mathbf{u}}, \hat{\mathbf{v}}$ be the empirical singular value and unit-norm singular vectors of \mathbf{X} . Then the oracle shrinkage factor is $\hat{c} = \arg \min_c \|c \cdot \hat{s} \hat{\mathbf{u}} \hat{\mathbf{v}}^{\top} - s\mathbf{u}\mathbf{v}^{\top}\|_F^2 = \frac{s}{\hat{s}} \langle \hat{\mathbf{u}}, \mathbf{u} \rangle \langle \hat{\mathbf{v}}, \mathbf{v} \rangle$.

Suppose we have also $\tilde{\mathbf{X}} = s\mathbf{u}\mathbf{v}^{\top} + \tilde{\epsilon}$ with the same signal component and independent noise. Let $\tilde{s} = \|\tilde{\mathbf{X}} \hat{\mathbf{v}}\|$ and $\tilde{\mathbf{u}} = \tilde{\mathbf{X}} \hat{\mathbf{v}} / \tilde{s}$ be the signal strength and factor scores for $\hat{\mathbf{v}}$. Then

$$\langle \tilde{\mathbf{u}}, \hat{\mathbf{u}} \rangle = \frac{1}{\tilde{s}} \langle (s\mathbf{u}\mathbf{v}^{\top} + \tilde{\epsilon}) \hat{\mathbf{v}}, \hat{\mathbf{u}} \rangle = \frac{s}{\tilde{s}} \langle \hat{\mathbf{v}}, \mathbf{v} \rangle \langle \hat{\mathbf{u}}, \mathbf{u} \rangle + \frac{1}{\tilde{s}} \hat{\mathbf{u}}^{\top} \tilde{\epsilon} \hat{\mathbf{v}} \approx \frac{s}{\tilde{s}} \langle \hat{\mathbf{v}}, \mathbf{v} \rangle \langle \hat{\mathbf{u}}, \mathbf{u} \rangle,$$

where the last approximation holds by independence of ϵ and $\tilde{\epsilon}$, as long as the entrywise noise variance of $\tilde{\epsilon}$ is much smaller than \tilde{s}^2 . Thus, $\frac{\tilde{s}}{s} \langle \tilde{\mathbf{u}}, \hat{\mathbf{u}} \rangle \approx \frac{s}{\tilde{s}} \langle \hat{\mathbf{v}}, \mathbf{v} \rangle \langle \hat{\mathbf{u}}, \mathbf{u} \rangle = \hat{c}$, so we may estimate the optimal shrinkage factor \hat{c} by $\frac{\tilde{s}}{s} \langle \tilde{\mathbf{u}}, \hat{\mathbf{u}} \rangle$.

Returning to SPACO, identifying $(\|\boldsymbol{\mu}_k\|_2, \frac{\boldsymbol{\mu}_k}{\|\boldsymbol{\mu}_k\|_2})$ with $(\hat{s}_k, \hat{\mathbf{u}}_k)$ of this example, this suggests to set

$$(11) \quad c_k = \frac{\|\tilde{\boldsymbol{\mu}}_k\|_2}{\|\boldsymbol{\mu}_k\|_2} \left(\langle \tilde{\boldsymbol{\mu}}_k, \boldsymbol{\mu}_k \rangle \right)_+$$

if we had access to $\tilde{\boldsymbol{\mu}}_k$ estimated from an independent copy of the signal array. For SPACO-without covariates, a natural correction based on this idea is to use cross-fitting, where we approximate $\tilde{\boldsymbol{\mu}}_k$ using model parameters estimated from data excluding different folds:

1. Divide the data into M different folds $\cup_{m=1}^M \mathcal{V}_m$.
2. Compute the posterior mean $\tilde{\boldsymbol{\mu}}$ of \mathbf{U} for subjects in fold \mathcal{V}_m using the parameters $\mathbf{V}^{-m}, \boldsymbol{\Phi}^{-m}$ estimated from data excluding fold \mathcal{V}_m .
3. For each k^{th} component, set the shrinkage factor c_k by eq. (11).

For SPACO with covariates \mathbf{Z} , the correction is more complicated because we want to correct potential inflation in $\mathbf{Z}\boldsymbol{\beta}$. For this purpose, we use a cross-fitted $\hat{\boldsymbol{\beta}}$. More specifically, in the $\boldsymbol{\beta}_k$ -update step of Lemma 3.1, we construct \tilde{y}_i for $i \in \mathcal{V}_m$ as

$$\tilde{y}_i = \sqrt{\frac{1}{s_k^2 s_k^2 - (\Sigma_i)_{k,k}}} \left((\Sigma_i)_{k,:} (\mathbf{V}^{-m} \odot \boldsymbol{\Phi}^{-m})^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} + (\Sigma_i)_{k,k^c} \Lambda_{f,k^c} \hat{\boldsymbol{\beta}}_{k^c}^\top \mathbf{z}_i \right)$$

in place of (10). In our empirical study, cross-fitting reduces the reconstruction error in high dimensions and offers better empirical type I error control when testing for variable importance.

4.4. Variable importance measure. We are often interested in measuring the importance of each covariate \mathbf{Z}_j on a given factor \mathbf{u}_k . In this section, we construct variable importance measures as approximate p-values from tests of conditional association or marginal association between \mathbf{Z}_j and \mathbf{u}_k . Both measures are popular in practice.

In Section 4.4.1, we first briefly describe hypothesis testing based on randomization, and its application to our problem when all model parameters other than $\boldsymbol{\beta}_k$ for the k^{th} factor are known. We then discuss testing with the estimated model parameters in Section 4.4.2.

4.4.1. Randomization test for SPACO with known model parameters. First consider a simple linear regression model

$$Y = \mathbf{Z}^\top \boldsymbol{\beta} + \zeta,$$

with mean-zero noise ζ independent of $\mathbf{Z} = (Z_1, \dots, Z_q)$. P-values for testing the conditional independence/partial correlation or the marginal independence/association

$$H_0^{\text{partial}} : \mathbf{Z}_j \perp\!\!\!\perp Y | \mathbf{Z}_{j^c}, \quad H_0^{\text{marginal}} : \mathbf{Z}_j \perp\!\!\!\perp Y$$

are popular measures of the importance of a feature \mathbf{Z}_j . Randomization test is a procedure for constructing p-values without assuming the correctness of the linear model of Y on \mathbf{Z} .

Suppose that we have access to the conditional distribution of $\mathbf{Z}_j | \mathbf{Z}_{j^c}$. Let $t(\mathbf{Z}_j, \mathbf{Z}_{j^c}, \mathbf{y})$ be a test statistic. Let $T := t(\mathbf{Z}_j, \mathbf{Z}_{j^c}, \mathbf{y})$ and $T^* := t(\mathbf{Z}_j^*, \mathbf{Z}_{j^c}, \mathbf{y})$, where \mathbf{Z}_j^* is an independent copy generated from the conditional distribution $\mathbf{Z}_j | \mathbf{Z}_{j^c}$ and \mathbf{Z}_{j^c} . Then, under the null hypothesis H_0^{partial} of conditional independence, T and T^* have the same law conditional on \mathbf{Z}_{j^c} and \mathbf{y} . So $P(T > t_{1-\alpha}^* | \mathbf{Z}_{j^c}, \mathbf{y}) \leq \alpha$ for any $\alpha \in (0, 1)$, where $t_{1-\alpha}^*$ is the

$(1 - \alpha)$ -percentile of the conditional distribution of T^* (Candès et al., 2018). Similarly, for the null hypothesis H_0^{marginal} of marginal independence, consider a test statistic $t(\mathbf{Z}_j, \mathbf{y})$ and set $T := t(\mathbf{Z}_j, \mathbf{y})$ and $T^* := t(\mathbf{Z}_j^*, \mathbf{y})$ where \mathbf{Z}_j^* is an independent copy generated from the marginal distribution of Z_j . Then $P(T > t_{1-\alpha}^* | \mathbf{y}) \leq \alpha$ for any $\alpha \in (0, 1)$.

For SPACO, we are interested in testing the analogous hypotheses

$$H_{0k}^{\text{partial}} : Z_j \perp\!\!\!\perp \mu_k | Z_{j^c}, \quad H_{0k}^{\text{marginal}} : Z_j \perp\!\!\!\perp \mu_k$$

for the relationship between \mathbf{Z}_j and a single factor μ_k . In an idealized setting where all model parameters except β_k are known, let $\mathbf{V}, \Phi, s^2, \sigma^2$ and β_ℓ^* for $\ell \neq k$ be the true model parameters. Then the SPACO estimate for β_k is the solution to a penalized regression problem:

PROPOSITION 4.2. *Given the other model parameters $\mathbf{V}, \Phi, s^2, \sigma^2$ and β_ℓ^* for $\ell \neq k$, the SPACO estimate for β_k is*

$$\hat{\beta}_k = \arg \min_{\beta_k} \left\{ \sum_{i=1}^I \frac{1}{w_i} (\tilde{\mathbf{z}}_i^\top \beta_k - \tilde{y}_i)^2 + \lambda_{2k} |\beta_k|_1 \right\} = \arg \min_{\beta_k} \left\{ \sum_{i=1}^I \frac{1}{w_i} (\tilde{\mathbf{z}}_i^\top \beta_k - \tilde{\mathbf{z}}_i^\top \beta_k^* - \zeta_i)^2 + \lambda_{2k} |\beta_k|_1 \right\},$$

where $w_i = s_k^2 - (\Sigma_i)_{kk}$, $\zeta_i \sim \mathcal{N}(0, w_i)$, and $\tilde{\mathbf{z}}_i, \tilde{y}_i$ are defined by eq. (10) at the given model parameters.

Hence, the testing problem reduces to that of testing the variable triple $(\tilde{Z}_j, \tilde{Z}_{j^c}, \tilde{Y})$ in a linear model, and we may perform the above randomization tests:

- For the partial correlation test, we use a residualized and weighted univariate regression coefficient

$$T = \frac{\sum_{i=1}^I \frac{1}{w_i} (\tilde{y}_i - \tilde{\mathbf{z}}_{i,j^c}^\top \hat{\beta}_{j^c,k}) \tilde{\mathbf{z}}_{ij}}{\sum_{i=1}^I \frac{1}{w_i} \tilde{\mathbf{z}}_{ij}^2}, \quad T_* = \frac{\sum_{i=1}^I \frac{1}{w_i} (\tilde{y}_i - \tilde{\mathbf{z}}_{i,j^c}^\top \hat{\beta}_{j^c,k}) \tilde{\mathbf{z}}_{ij}^*}{\sum_{i=1}^I \frac{1}{w_i} \tilde{\mathbf{z}}_{ij}^{*2}}.$$

Here $\hat{\beta}_{j^c,k}$ is trained using the weighted regression in Proposition 4.2 excluding covariate j , and $\tilde{\mathbf{z}}_{ij}^*$ is constructed as in eq. (9) replacing \mathbf{z}_{ij} with \mathbf{z}_{ij}^* , which are independently generated from the conditional distribution of $Z_j | Z_{j^c}$.

- For the marginal correlation test, we use the weighted univariate regression coefficient

$$T = \frac{\sum_{i=1}^I \frac{1}{w_i} \tilde{y}_i \tilde{\mathbf{z}}_{ij}}{\sum_{i=1}^I \frac{1}{w_i} \tilde{\mathbf{z}}_{ij}^2}, \quad T_* = \frac{\sum_{i=1}^I \frac{1}{w_i} \tilde{y}_i \tilde{\mathbf{z}}_{ij}^*}{\sum_{i=1}^I \frac{1}{w_i} \tilde{\mathbf{z}}_{ij}^{*2}},$$

with $\tilde{\mathbf{z}}_{ij}^*$ constructed as in eq. (9) replacing \mathbf{z}_{ij} with \mathbf{z}_{ij}^* independently generated from the marginal distribution of Z_j .

4.4.2. *Testing for SPACO with estimated parameters.* In practice, β_ℓ for the other factors $\ell \neq k$ are also estimated from the data. Inaccurate estimation of β_ℓ^* may directly impact the validity of the above randomization test for β_k . In this section, we provide a valid test by constructing alternative response variables, covariates, and weights that do not depend on β_ℓ for $\ell \neq k$.

More generally, we consider a family of response and covariate constructions indexed by $\delta \in [0, 1]$, where $\delta = 1$ corresponds to the original update for β_k in SPACO, and $\delta = 0$ fully decouples β_k from β_ℓ .

LEMMA 4.3. Fixing the model parameters $\mathbf{V}, \Phi, s^2, \sigma^2$, let β_ℓ be an estimate of β_ℓ^* for $\ell \neq k$. Define $\Sigma_i(\delta) = \left(\delta \Lambda_f + (\mathbf{V} \odot \Phi)_{i,\cdot}^\top \Lambda_{\tilde{i}}^{-1} (\mathbf{V} \odot \Phi)_{i,\cdot} \right)^{-1}$,

$$\tilde{\mathbf{z}}_i(\delta) = \left(1 - \delta \frac{\Sigma_i(\delta)_{kk}}{s_k^2} \right) \mathbf{z}_i,$$

$$(12) \quad \tilde{y}_i(\delta) = \Sigma_i(\delta)_{k,\cdot} (\mathbf{V} \odot \Phi)_{i,\cdot}^\top \Lambda_{\tilde{i}}^{-1} \mathbf{X}_{I,\tilde{i}} + \delta \sum_{\ell \neq k} \frac{\Sigma_i(\delta)_{k\ell}}{s_\ell^2} \mathbf{z}_i^\top \beta_\ell,$$

and $w_i(\delta) = s_k^2 + (1 - 2\delta)\Sigma_i(\delta)_{kk} + (\delta^2 - \delta)\Sigma_i(\delta)_{k,\cdot} \Lambda_f \Sigma_i(\delta)_{\cdot,k}$. Then $\tilde{y}_i(\delta)$ has the distribution

$$\tilde{y}_i(\delta) = \tilde{\mathbf{z}}_i(\delta)^\top \beta_k^* + \delta \sum_{\ell \neq k} \frac{\Sigma_i(\delta)_{k\ell}}{s_\ell^2} \mathbf{z}_i^\top (\beta_\ell - \beta_\ell^*) + \xi_i, \quad \xi_i \sim \mathcal{N}(0, w_i(\delta)).$$

The above dependence of $\tilde{y}_i(\delta)$ on $\mathbf{z}_i^\top (\beta_\ell - \beta_\ell^*)$ may render the randomization test invalid when $\beta_\ell \neq \beta_\ell^*$. Thus, we typically set $\delta = 0$ to remove this dependence. Under this setting, we have $\tilde{\mathbf{z}}_i(0) = \mathbf{z}_i$ and $\tilde{y}_i(0) = \mathbf{z}_i^\top \beta_k^* + \xi_i$. Algorithm 3 summarizes our proposal for performing hypothesis testing using these constructed response variables and weights.

Algorithm 3: Randomization test for \mathbf{Z}_j

```

1 for  $k = 1, \dots, K$  do
3   Construct responses and features as described in Lemma 4.3, for  $\delta = 0$ .
5   Define  $\hat{\beta}_k$  by
      
$$\hat{\beta}_k = \arg \min_{\beta_k: \beta_{k,j}=0} \left\{ \sum_{i=1}^I \frac{1}{w_i(0)} (\mathbf{z}_i^\top \beta_k - \tilde{y}_i(0))^2 + \lambda_{2k} |\beta_k|_1 \right\}.$$

7   Compute the designed test statistic  $T$  using  $(\mathbf{Z}_j, \mathbf{Z}_{j^c}, \tilde{\mathbf{y}}(0), \hat{\beta}_{j^c,k})$ .
9   Compute randomized statistics  $T_b^*$  using  $(\mathbf{Z}_j^{b*}, \mathbf{Z}_{j^c}, \tilde{\mathbf{y}}(0), \hat{\beta}_{j^c,k})$ , where  $\mathbf{Z}_j^{b*}$  for  $b = 1, \dots, B$  are the generated (conditionally or marginally) independent copies of  $\mathbf{Z}_j$ .
11  Let  $\hat{G}(\cdot)$  be the empirical estimate of the CDF of  $T$  under  $H_0$  using  $\{T_1^*, \dots, T_B^*\}$ , and return the two-sided p-value  $p = [1 - \hat{G}(|T|)] + \hat{G}(-|T|)$ .
12 end
```

LEMMA 4.4. Suppose the model parameters \mathbf{V}, Φ, s^2 and σ^2 are known. Fix j, k , and let $\tilde{y}_i(0)$ and $w(0)$ be the constructed responses and weights in Lemma 4.3 with $\delta = 0$. Let $\hat{\beta}_k$ be the estimate in Algorithm 3, and define

$$T_{\text{partial}} = \frac{\sum_i \frac{1}{w_i(0)} (\tilde{y}_i(0) - \mathbf{z}_{i,j^c}^\top \hat{\beta}_{j^c,k}) z_{ij}}{\sum_i \frac{1}{w_i(0)} z_{ij}^2}, \quad T_{\text{marginal}} = \frac{\sum_i \frac{1}{w_i(0)} \tilde{y}_i z_{ij}}{\sum_i \frac{1}{w_i(0)} z_{ij}^2}.$$

Let \mathbf{Z}_j^* be the generated (conditionally or marginally) independent copy of \mathbf{Z}_j , and define $T_{\text{partial}}^*, T_{\text{marginal}}^*$ similarly with \mathbf{Z}_j^* in place of \mathbf{Z}_j . Let $\tilde{\mathbf{y}}(0) = (\tilde{y}_1(0), \dots, \tilde{y}_I(0))^T$, then

- $T_{\text{partial}} | (\tilde{\mathbf{y}}(0), \mathbf{Z}_{j^c}) \stackrel{d}{=} T_{\text{partial}}^* | (\tilde{\mathbf{y}}(0), \mathbf{Z}_{j^c})$ if H_{0k}^{partial} holds.
- $T_{\text{marginal}} | \tilde{\mathbf{y}}(0) \stackrel{d}{=} T_{\text{marginal}}^* | \tilde{\mathbf{y}}(0)$ if H_{0k}^{marginal} holds.

REMARK 4.4.1. *More generally, for $\delta > 0$, the test based on $\tilde{y}_i(\delta)$, $w(\delta)$, and $\tilde{\mathbf{z}}_i(\delta)$ trades off robustness against estimation errors in $Z^\top(\boldsymbol{\beta}_\ell - \boldsymbol{\beta}_\ell^*)$ for possibly increased power.*

To see this, suppose that $\boldsymbol{\beta}_\ell = \boldsymbol{\beta}_\ell^$ for $\ell \neq k$ and $Z_j \sim \mathcal{N}(0, 1)$. The signal to noise ratio with $\delta \in \{0, 1\}$ can be calculated as*

$$\begin{aligned} \text{SNR}(0) &= \frac{\mathbb{E}(\mathbf{z}_i^\top \boldsymbol{\beta}_k^*)^2}{\frac{1}{I} \sum_i w_i(0)} = \frac{\|\boldsymbol{\beta}_k^*\|_2^2}{\frac{1}{I} \sum_i w_i(0)} = \frac{\|\boldsymbol{\beta}_k^*\|_2^2}{s_k^2 + \frac{1}{I} \sum_i \Sigma_i(0)_{kk}}, \\ \text{SNR}(1) &= \frac{\frac{1}{I} \sum_i \mathbb{E}(\tilde{\mathbf{z}}_i(1)^\top \boldsymbol{\beta}_k^*)^2}{\frac{1}{I} \sum_i w_i(1)} = \frac{\frac{1}{I} \sum_i (1 - \frac{\Sigma_i(1)_{kk}}{s_k^2}) \|\boldsymbol{\beta}_k^*\|_2^2}{s_k^2 - \frac{1}{I} \sum_i \Sigma_i(1)_{kk}} = \frac{\|\boldsymbol{\beta}_k^*\|_2^2}{s_k^2}. \end{aligned}$$

Thus, the signal-to-noise ratio is higher with $\delta = 1$ if we have access to the true $\boldsymbol{\beta}_\ell^$.*

In practice, the model parameters \mathbf{V} , Φ , s^2 , σ^2 of SPACO are also unknown, and we will substitute their empirical estimates in the above procedure. In particular, we use \mathbf{V} , Φ from the cross-fitted SPACO for constructing $\tilde{y}(0)$, i.e., as described in Section 4.3, for $i \in \mathcal{V}_m$ we construct

$$\tilde{y}_i(0) = \Sigma_i(0)_{k,:} (\mathbf{V}^{-m} \odot \Phi^{-m})^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}}.$$

The randomization test requires generating Z_j^* from the conditional or marginal distribution of Z_j , and estimating the resulting distribution of T^* . We will estimate the distribution of T^* by fitting a skewed t-distribution as suggested in (Katsevich and Roeder, 2020). The use of the fitted $\hat{G}(\cdot)$ instead of the empirical CDF can greatly reduce the computational cost: We may obtain accurate estimates of small p-values around 10^{-4} using only 100 independent generations of Z_j^* and the fitted $\hat{G}(\cdot)$. More details on the generations of Z_j^* are provided in Appendix B.3.

5. Simulations.

5.1. *Comparison between SPACO and SPACO-*. In this section, we compare SPACO with SPACO- (which does not use the covariates \mathbf{Z}) for reconstruction of the underlying signal array, as measured by Pearson correlation between the estimated values and the signal values. We simulate Gaussian data across different signal levels and different missing rates along the time dimension:

- We test $(I, T, J, q) = (100, 30, 10, 100)$ and $(100, 30, 500, 100)$. The number of factors is $K = 3$, the noise variance of ϵ is fixed at 1, and \mathbf{V} and \mathbf{Z} are randomly generated as $V_{jk} \sim \mathcal{N}(0, \frac{1}{J})$ and $Z_{i\ell} \sim \mathcal{N}(0, 1)$.
- We observe a $\gamma = \{100\%, 50\%, 10\%\}$ fraction of time points $\{1, \dots, T\}$, chosen independently for each subject.
- We set $\phi_1(t) = \theta_1$, $\phi_2(t) = \theta_2 \sqrt{1 - (\frac{t}{T})^2}$, $\phi_3(t) = \theta_3 \cos(4\pi \frac{t}{T})$ with random parameters $\theta_1, \theta_2, \theta_3 \sim c_1 \cdot \mathcal{N}(0, \frac{\log J + \log T}{nT\gamma})$ for $c_1 \in \{1, 3, 5\}$.
- We set $\beta_{\ell,k} \sim c_2 \cdot \mathcal{N}(0, \frac{\log q}{I})$ for $c_2 \in \{0, 3, 10\}$ for the first 3 covariates $\ell = 1, 2, 3$, and $\beta_{\ell,k} = 0$ otherwise. When $c_2 = 0$, all $\beta_{\ell,k}$ are 0. Each U_k is standardized to have mean 0 and variance 1 after generation.

Consistent with our expectation, SPACO improves upon SPACO- when \mathbf{X} contains insufficient information to determine the underlying truth while \mathbf{Z} is informative. Table 2 shows the average correlations between the reconstructed and true signals cross 50 random repetitions. In these simulations, SPACO is comparable to SPACO- when the signal in \mathbf{X} is strong or when the signal in \mathbf{Z} is weak; it provides a significant gain in settings that favor SPACO.

TABLE 2

Correlations between the reconstructed values and the underlying signals: raw = raw observation \mathbf{X} . For each simulation set-up, we use boldface to indicate methods that are within 2 standard deviations from the top performer. r in the column names represents percent of observed time points per subject. SNR represents the signal strength c_2 for $\beta_{\ell,k}$, e.g., $r1SNR3$ corresponds to observation rate $\gamma = 10\%$ and $c_2 = 3$. Similarly, SNR- and J - in the row names represent the values for c_1 and J .

setting	method	r1SNR0	r1SNR3	r1SNR10	r.5SNR0	r.5SNR3	r.5SNR10	r.1SNR0	r.1SNR3	r.1SNR10
SNR1J10	raw	0.059(1E-03)	0.057(1E-03)	0.057(1E-03)	0.078(2E-03)	0.078(2E-03)	0.083(2E-03)	0.186(4E-03)	0.177(3E-03)	0.178(4E-03)
	SPACO-	0.156(1E-02)	0.136(1E-02)	0.165(2E-02)	0.107(1E-02)	0.145(1E-02)	0.172(2E-02)	0.198(1E-02)	0.19(1E-02)	0.208(1E-02)
	SPACO	0.154(1E-02)	0.167(2E-02)	0.26(2E-02)	0.101(9E-03)	0.174(2E-02)	0.269(2E-02)	0.182(1E-02)	0.208(1E-02)	0.298(2E-02)
SNR3J10	raw	0.169(3E-03)	0.169(3E-03)	0.168(3E-03)	0.23(4E-03)	0.228(5E-03)	0.24(5E-03)	0.466(7E-03)	0.475(7E-03)	0.484(7E-03)
	SPACO-	0.786(1E-02)	0.796(1E-02)	0.789(1E-02)	0.795(9E-03)	0.778(1E-02)	0.804(9E-03)	0.793(9E-03)	0.802(9E-03)	0.811(9E-03)
	SPACO	0.787(1E-02)	0.818(9E-03)	0.859(8E-03)	0.793(9E-03)	0.798(1E-02)	0.87(7E-03)	0.788(9E-03)	0.821(8E-03)	0.886(7E-03)
SNR5J10	raw	0.271(5E-03)	0.269(5E-03)	0.275(5E-03)	0.378(6E-03)	0.37(6E-03)	0.374(7E-03)	0.666(6E-03)	0.65(7E-03)	0.656(7E-03)
	SPACO-	0.924(4E-03)	0.924(3E-03)	0.92(6E-03)	0.926(3E-03)	0.926(3E-03)	0.923(4E-03)	0.921(3E-03)	0.9(5E-03)	0.915(3E-03)
	SPACO	0.925(3E-03)	0.929(3E-03)	0.944(6E-03)	0.927(3E-03)	0.93(3E-03)	0.946(3E-03)	0.92(3E-03)	0.9(5E-03)	0.944(2E-03)
SNR1J500	raw	0.014(1E-04)	0.014(1E-04)	0.014(1E-04)	0.019(2E-04)	0.019(2E-04)	0.019(2E-04)	0.043(4E-04)	0.043(4E-04)	0.043(4E-04)
	SPACO-	0.012(1E-03)	0.014(2E-03)	0.013(2E-03)	0.011(1E-03)	0.011(1E-03)	0.016(2E-03)	0.015(1E-03)	0.016(1E-03)	0.018(2E-03)
	spaco	0.012(1E-03)	0.019(2E-03)	0.019(3E-03)	0.011(1E-03)	0.012(1E-03)	0.018(2E-03)	0.015(1E-03)	0.017(2E-03)	0.021(3E-03)
SNR3J500	raw	0.041(2E-04)	0.041(2E-04)	0.041(2E-04)	0.057(2E-04)	0.058(2E-04)	0.058(2E-04)	0.128(6E-04)	0.128(5E-04)	0.129(6E-04)
	SPACO-	0.451(1E-02)	0.476(1E-02)	0.487(1E-02)	0.529(2E-02)	0.547(1E-02)	0.547(1E-02)	0.665(6E-03)	0.649(7E-03)	0.615(1E-02)
	SPACO	0.451(1E-02)	0.491(2E-02)	0.513(1E-02)	0.53(2E-02)	0.555(2E-02)	0.583(1E-02)	0.665(6E-03)	0.662(7E-03)	0.661(1E-02)
SNR5J500	raw	0.068(2E-04)	0.068(2E-04)	0.068(2E-04)	0.096(2E-04)	0.096(3E-04)	0.095(4E-04)	0.21(9E-04)	0.212(7E-04)	0.21(8E-04)
	SPACO-	0.875(1E-03)	0.865(4E-03)	0.856(4E-03)	0.875(6E-04)	0.874(1E-03)	0.861(3E-03)	0.872(1E-03)	0.873(1E-03)	0.867(2E-03)
	SPACO	0.874(1E-03)	0.866(4E-03)	0.863(4E-03)	0.875(6E-04)	0.875(1E-03)	0.869(3E-03)	0.872(1E-03)	0.874(9E-04)	0.877(1E-03)

TABLE 3

Relative MSE between cross-fitted SPACO/SPACO- and SPACO/SPACO-: Each value is the ratio between MSE of the cross-fitted procedure to the procedure without cross-fitting. The column labels r - and SNR- and row labels SNR- and J - are the same as in Table 2.

setting	method	r1SNR0	r1SNR3	r1SNR10	r.5SNR0	r.5SNR3	r.5SNR10	r.1SNR0	r.1SNR3	r.1SNR10
SNR1J10	spaco0	0.906	0.9	0.901	0.905	0.915	0.917	0.942	0.945	0.933
	spaco	0.863	0.847	0.843	0.873	0.879	0.854	0.862	0.879	0.84
SNR3J10	spaco0	0.983	0.982	0.979	0.998	0.994	0.991	1.02	1.02	1.02
	spaco	0.982	0.989	0.974	1	1	0.983	1.03	1.04	1.01
SNR5J10	spaco0	0.998	1	0.998	1.02	1.02	1.02	1.09	1.05	1.12
	spaco	1	1.02	1.04	1.02	1.05	1.05	1.09	1.07	1.11
SNR1J500	spaco0	0.26	0.262	0.259	0.263	0.267	0.268	0.326	0.33	0.296
	spaco	0.242	0.242	0.244	0.256	0.254	0.26	0.337	0.307	0.302
SNR3J500	spaco0	0.628	0.62	0.595	0.639	0.624	0.601	0.701	0.708	0.677
	spaco	0.628	0.627	0.598	0.64	0.631	0.601	0.696	0.706	0.679
SNR5J500	spaco0	0.841	0.826	0.784	0.866	0.861	0.813	0.926	0.926	0.908
	spaco	0.843	0.836	0.815	0.867	0.867	0.835	0.927	0.932	0.937

5.2. *Improvements using cross-fitting.* In reconstructing the signal array, SPACO/SPACO- can overfit towards the noise, resulting in high reconstruction errors in weak signal and high-dimensional settings. Here, we compare the reconstruction errors with and without cross-fitting in the same simulation set-ups as in Section 5.1. In each simulation set-up, we treat MSE for SPACO-/SPACO as the baseline and show the relative levels of errors after cross-fitting. The corrected SPACO/SPACO- estimates achieve much lower reconstruction errors in high-dimensional and weak-signal settings, with a slight increase in error in the low-dimensional and strong-signal settings.

TABLE 4

Conditional association test: Achieved type-I error and power with target level $\alpha = 0.01, 0.05$ and simulations described in Section 5.1, averaged across 50 repetitions. For the rows, the first column of row names SNR- J represent the values of c_1 and J , and the second column of row names SNR- r represent c_2 and the observation rate γ . All associations are null when $c_2 = 0$, and some are null when $c_2 = 10$.

Xsignal	Zsignal+sparsity	0.01				0.05			
		SPACO-typeI	power	SPACO-typeI	power	SPACO-typeI	power	SPACO-typeI	power
SNR1J10	SNR0r1.0	0.012	nan	0.01	nan	0.054	nan	0.053	nan
	SNR0r0.5	0.012	nan	0.01	nan	0.055	nan	0.057	nan
	SNR0r0.1	0.011	nan	0.01	nan	0.058	nan	0.057	nan
	SNR10r1.0	0.012	0.15	0.011	0.14	0.054	0.24	0.052	0.22
	SNR10r0.5	0.011	0.16	0.0093	0.12	0.053	0.25	0.053	0.23
	SNR10r0.1	0.013	0.14	0.011	0.12	0.058	0.23	0.056	0.2
SNR3J10	SNR0r1.0	0.0096	nan	0.01	nan	0.053	nan	0.054	nan
	SNR0r0.5	0.0083	nan	0.0094	nan	0.053	nan	0.053	nan
	SNR0r0.1	0.011	nan	0.011	nan	0.051	nan	0.052	nan
	SNR10r1.0	0.0096	0.54	0.0093	0.54	0.052	0.61	0.055	0.64
	SNR10r0.5	0.0096	0.59	0.01	0.6	0.054	0.64	0.054	0.66
	SNR10r0.1	0.0096	0.58	0.011	0.6	0.052	0.67	0.053	0.69
SNR1J500	SNR0r1.0	0.016	nan	0.012	nan	0.059	nan	0.058	nan
	SNR0r0.5	0.012	nan	0.01	nan	0.054	nan	0.051	nan
	SNR0r0.1	0.0088	nan	0.0096	nan	0.048	nan	0.05	nan
	SNR10r1.0	0.013	0.04	0.01	0.038	0.058	0.12	0.054	0.1
	SNR10r0.5	0.011	0.047	0.012	0.044	0.053	0.13	0.054	0.12
	SNR10r0.1	0.0097	0.042	0.011	0.033	0.055	0.12	0.055	0.12
SNR3J500	SNR0r1.0	0.0094	nan	0.0087	nan	0.054	nan	0.054	nan
	SNR0r0.5	0.0085	nan	0.0095	nan	0.051	nan	0.053	nan
	SNR0r0.1	0.0091	nan	0.011	nan	0.048	nan	0.051	nan
	SNR10r1.0	0.013	0.48	0.012	0.48	0.057	0.57	0.056	0.58
	SNR10r0.5	0.01	0.54	0.012	0.55	0.054	0.61	0.057	0.61
	SNR10r0.1	0.0089	0.64	0.01	0.63	0.052	0.72	0.056	0.72

TABLE 5

Same as Table 4, for the marginal association test.

Xsignal	Zsignal+sparsity	0.01				0.05			
		SPACO-typeI	power	SPACO-typeI	power	SPACO-typeI	power	SPACO-typeI	power
SNR1J10	SNR0r1.0	0.011	nan	0.01	nan	0.056	nan	0.053	nan
	SNR0r0.5	0.012	nan	0.01	nan	0.057	nan	0.057	nan
	SNR0r0.1	0.013	nan	0.011	nan	0.059	nan	0.056	nan
	SNR10r1.0	0.013	0.15	0.011	0.14	0.055	0.24	0.051	0.22
	SNR10r0.5	0.012	0.16	0.0091	0.12	0.053	0.23	0.053	0.24
	SNR10r0.1	0.013	0.13	0.01	0.11	0.06	0.22	0.057	0.2
SNR3J10	SNR0r1.0	0.0097	nan	0.011	nan	0.054	nan	0.055	nan
	SNR0r0.5	0.0085	nan	0.0094	nan	0.052	nan	0.053	nan
	SNR0r0.1	0.011	nan	0.011	nan	0.051	nan	0.052	nan
	SNR10r1.0	0.01	0.5	0.011	0.5	0.053	0.58	0.054	0.59
	SNR10r0.5	0.01	0.53	0.0096	0.53	0.054	0.61	0.055	0.6
	SNR10r0.1	0.0094	0.54	0.011	0.56	0.053	0.62	0.054	0.64
SNR1J500	SNR0r1.0	0.015	nan	0.012	nan	0.066	nan	0.058	nan
	SNR0r0.5	0.011	nan	0.01	nan	0.054	nan	0.051	nan
	SNR0r0.1	0.0086	nan	0.0096	nan	0.048	nan	0.051	nan
	SNR10r1.0	0.015	0.04	0.011	0.038	0.063	0.13	0.054	0.11
	SNR10r0.5	0.011	0.042	0.011	0.044	0.053	0.13	0.053	0.12
	SNR10r0.1	0.0098	0.04	0.011	0.031	0.056	0.11	0.057	0.12
SNR3J500	SNR0r1.0	0.0097	nan	0.0091	nan	0.054	nan	0.054	nan
	SNR0r0.5	0.0083	nan	0.0094	nan	0.051	nan	0.054	nan
	SNR0r0.1	0.0091	nan	0.011	nan	0.048	nan	0.051	nan
	SNR10r1.0	0.014	0.42	0.012	0.42	0.057	0.52	0.059	0.54
	SNR10r0.5	0.011	0.46	0.01	0.47	0.056	0.56	0.055	0.56
	SNR10r0.1	0.011	0.58	0.011	0.57	0.055	0.65	0.057	0.66

5.3. Variable importance and p-values. We investigate the approximated p-values based on cross-fitting for testing both the partial and marginal associations of \mathbf{Z} with $\boldsymbol{\mu}$. Recall that the conditional association test uses conditional randomization and residuals, while the

marginal test uses unconditional randomization and \tilde{y} . Since our variables in \mathbf{Z} are generated independently, the two null hypotheses coincide in this setting. Note that both SPACO- and SPACO are testing association with the covariates \mathbf{Z} , although SPACO- does not use \mathbf{Z} in model fitting.

Table 4 and Table 5 show the achieved Type I error and power at levels $\alpha = 0.01, 0.05$. Both SPACO and SPACO- achieve reasonable empirical control of the Type I error. Cross-fitting is important for ensuring this Type I error control: In Appendix C, we present qq-plots comparing p-values using cross-fitted \mathbf{V} and Φ and those using the in-sample \mathbf{V} and Φ , across 16 simulation set-ups. We observe a large deviation from the uniform distribution for the later construction.

The conditional association test yields higher power over the marginal test in the stronger signal settings, because we reduce the noise level when using the residual. The power of cross-fitted SPACO is better than SPACO- when the signal in \mathbf{X} is weak. This advantage disappears as the task becomes easy for both methods.

6. Applications. In this section, we demonstrate SPACO with immunological data from two recent SARS-COV-2 studies.

6.1. *Deep immune profiling of COVID-19 patients.* Data set A is from [Rendeiro et al. \(2020\)](#). It contains 45 longitudinal observations for individuals with COVID-19. For each individual, blood samples were collected at enrollment, and approximately every seven days after that if permissible. The authors performed deep immune profiling for seven immune panels at those time points. In addition to the longitudinal immune profiling data, covariates such as patient severity group and age were also collected.

We retain a set of immune profiles that are non-missing at every available sample, and perform quantile normalization of these features. For each sample, we consider days-from-symptom-onsite (DFS0) as its measured time. The final \mathbf{X} is of dimension $(I, T, J) = (36, 40, 133)$. The covariates \mathbf{Z} are of dimension $q = 10$. For each feature, we have removed the mean curve whose smoothness level is tuned via leave-one-subject-out cross-validation (LOSOCV). The average number of time points observed per subject is 1.86. The left half of Figure 1 shows examples of time trajectories of four features for different subjects, after this mean removal. Each line/dot represents observations from a single subject, with a dot indicating that only 1 time point was observed. Subjects with/without intubation are colored red/blue.

We apply SPACO to dataset A with $K = 5$ factors. The right half of Figure 1 shows example plots of raw observations against our reconstructions across all samples. We see a positive correlation between these observed and reconstructed values.

We evaluate the overall predictive power of our model based on LOSOCV. For each subject i , our prediction for x_{itj} is

$$\hat{x}_{itj} = \sum_{k=1}^K \hat{u}_{ik} \phi_{kt} v_{jk}.$$

Table 6 shows the average percent of variance explained across subjects, based on in-sample measures and LOSOCV. There are two types of in-sample measures. The first column “in-sample” uses the posterior mean μ from SPACO for the above values \hat{u}_{ik} . The second and third columns “insample0” and “cv.insample0” use only the estimated β from SPACO, with β of the former estimated using lasso penalties $\lambda_{2k} = 0$ and β of the later estimated using lasso penalties chosen by cross-validation (fixing other model parameters). The columns “cv.SPACO” and “cv.SPACO-” are reconstruction errors using LOSOCV, where we use $\hat{\beta}^{-i}$,

FIG 1. *Dataset A. Left panel: examples of time trajectories of four features for different subjects, after mean removal. Each line/dot represents observations from a single subject. Subjects with intubation are colored red. Horizontal time axis is the DFSO (self-reported). Right panel: Plots of observed feature values against estimated ones from SPACO. Each line/dot represents observations from a single subject.*

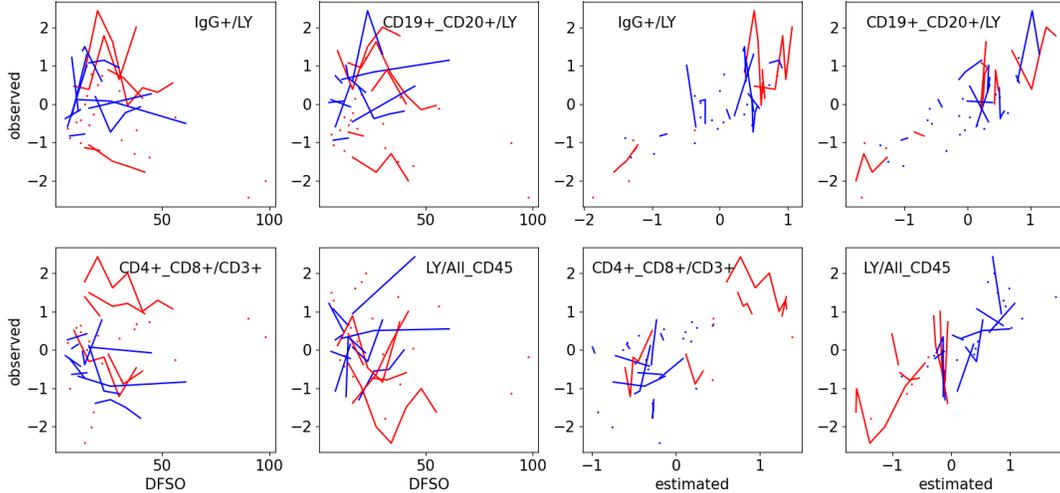


TABLE 6

Dataset A. Average percents of variance explained across subjects using in-sample data and LOSOCV.

insample	insample0	cv.insample0	cv.SPACO	cv.SPACO-
47.5	18.5	11.2	10.3	7.7

V^{-i} , and Φ^{-i} estimated excluding subject i . Comparing “cv_SPACO” and “cv_SPACO-”, “cv_SPACO-” estimates $\hat{\beta}^{-i}$ in a post-processing step after fitting the other model parameters. Figure 2 shows the negative logarithms of the p values from the randomized association tests, both for partial correlation and for marginal association.

6.2. *IMPACT COVID-19 study.* Data set B is from Lucas et al. (2020). It contains immune profiles for patients with COVID-19 from the IMPACT (Implementing Medical and Public Health Action Against Coronavirus CT) study. The authors assessed (1) levels of plasma cytokines and chemokines, (2) leukocyte profiles (by flow cytometry using freshly isolated peripheral blood mononuclear cells), and (3) PBMCs at various time points. Clinical variables including clinical scores, age are available.

We filter out features with more than 20% missingness among collected sample, and impute missing data using MOFA Argelaguet et al. (2018). The processed \mathbf{X} has a size $(I, T, J) = (98, 35, 111)$, and \mathbf{Z} is of dimension $q = 12$. The average number of time points observed per subject is 1.84. We consider a SPACO model with $K = 5$ factors, and perform the same analysis as we have done in dataset A. The left half of Figure 3 shows the examples of time trajectories of four features for different subjects after the mean removal, and the right half of Figure 3 plots the observed values against our estimated ones, with subjects in ICU colored red.

Table 7 shows the average percents of variance explained across subjects based on in-sample measures and LOSOCV. Figure 4 shows the negative log p-values from the randomized test.

FIG 2. Dataset A. P -values testing for partial correlation (left) and marginal association (right). The y -axis shows $-\log_{10}(\text{pvalue})$ and x -axis shows the covariate names. The horizontal blue line represents p -value being 0.05.

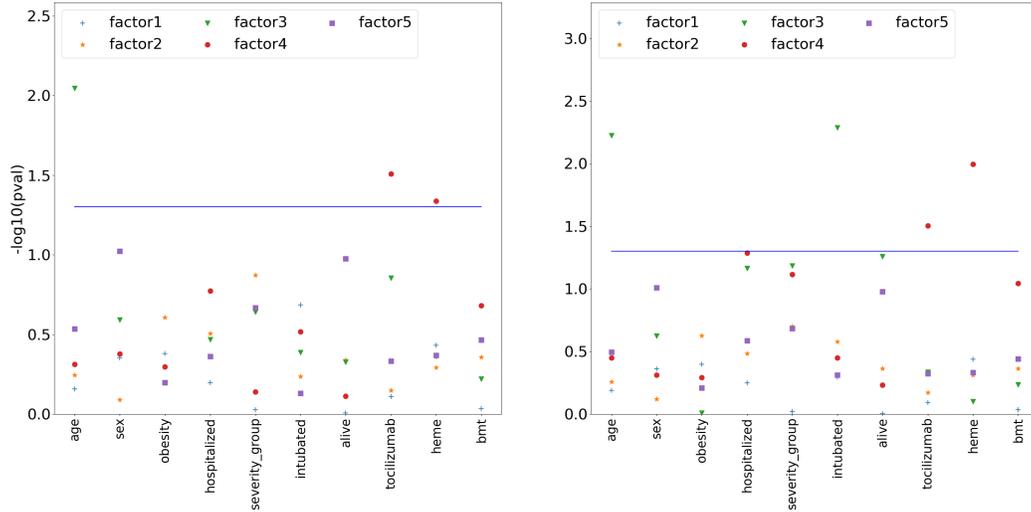


FIG 3. Dataset B. Left panel: examples of time trajectories of four features for different subjects, after mean removal. Subjects in ICU are colored red. Horizontal time axis is the self-reported DFSO. Right panel: Plots of observed and estimated feature values.

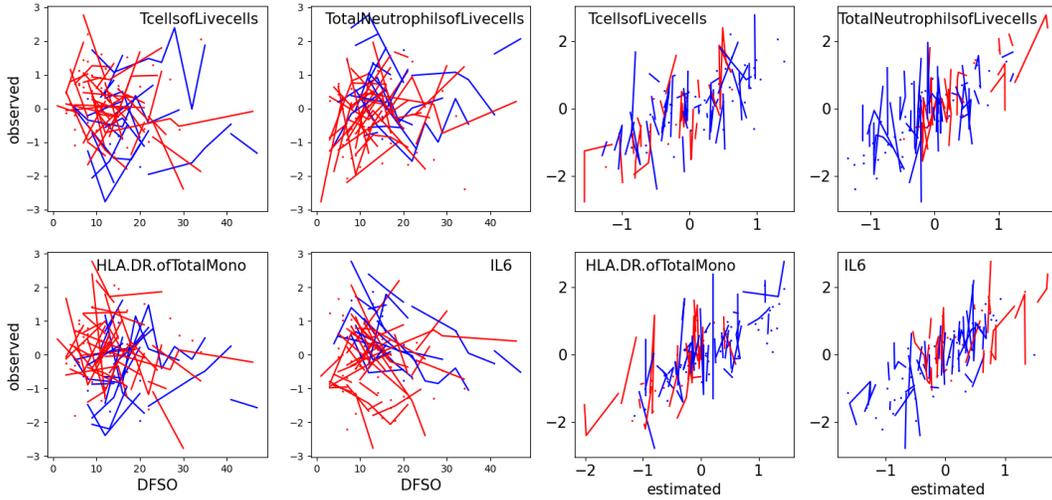
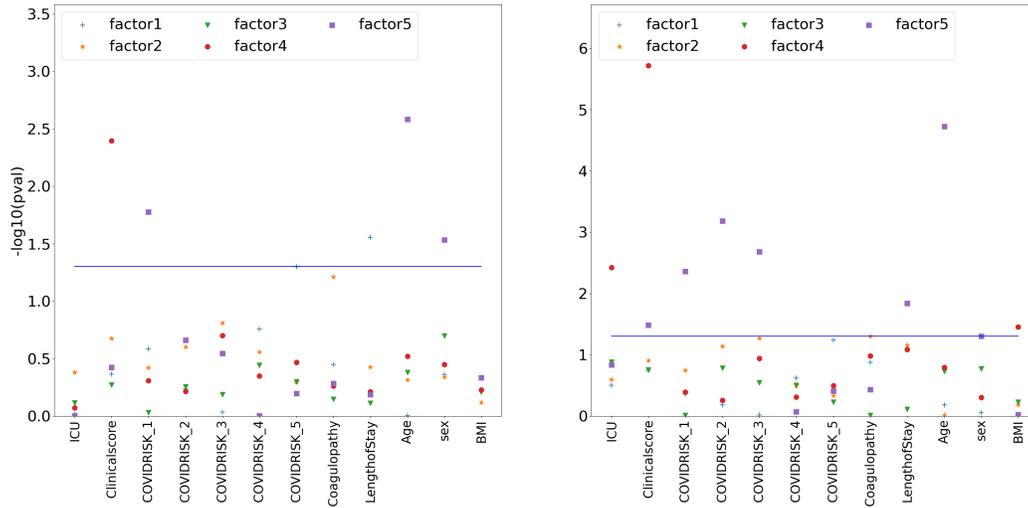


TABLE 7

Dataset B. Average percents of variance explained across subjects using in-sample data and LOSOCV.

insample	insample0	cv.insample0	cv.SPACO	cv.SPACO-
33.4	8.9	5.1	4.7	3.1

FIG 4. Dataset B. P -values testing for partial correlation (left) and marginal association (right). The y -axis shows $-\log_{10}(\text{pvalue})$ and x -axis shows the covariate names. The horizontal blue line represents p -value being 0.05.



6.3. *A closer look at the analysis results.* Several features in Z seem to be related to the derived factors in both datasets A and B. An age-related factor is identified in both data sets (factor 5). “Clinicalscores” is important in dataset B: It is associated with factor 5, and also associated strongly with factor 4, whose association with age is not strong.

Apart from claims about significance and reconstructions, we can also look into the composition of factors related to the clinical covariates. We take dataset B as an example, and demonstrate how other information from SPACO may be interesting to immunologists.

We can evaluate the contribution of features to each of the components by examining the corresponding entries in V . Figure 5 shows the top 25 features measured by the absolute value of V for factors 4 and 5. We divide subjects by age into young (≤ 70) and old (> 70), and by clinical score into mild (≤ 3) and severe (> 3), and we visualize these two covariates’ marginal effects on factors 4 and 5. To do this, we plot the weighted time curves for each of the participant groups, where we weight each of the features by their corresponding entries from V , i.e., $y_{it} = \mathbf{x}_{i,t}^\top V_k$ for factor k . Figure 6 shows the time trajectories for different patient groups. It seems that different disease severity groups start to diverge earlier using a feature combination specified by factor 4, while different age groups diverge later using a feature combination determined by factor 5. When combining the feature contribution coefficients V and the factor trajectory, we can qualitatively characterize the original feature’s behavior and cross-check with findings from the literature. For instance, the top contributor for factor 4 is IL-6. It is the most informative feature for separating the severe group from the mild group in Lucas et al. (2020), with severe patients showing higher IL-6 levels.

7. Discussion. We propose a model called SPACO that can jointly model sparse multivariate longitudinal data and auxiliary covariates. In our simulations, access to informative auxiliary covariates can improve the time series data reconstruction. We applied the proposed method to COVID-19 data and identified clinical variables that explain the heterogeneity between different patient subgroups. Even though both data sets are highly sparse and of small size, SPACO finds feature groups and time trajectories that can be potentially interesting for further investigation.

FIG 5. Bar plots of feature contribution in the construction of factors 4 and 5, measured by entries in V . The x-axis shows the corresponding values from V for different features that are indicated by the row names.

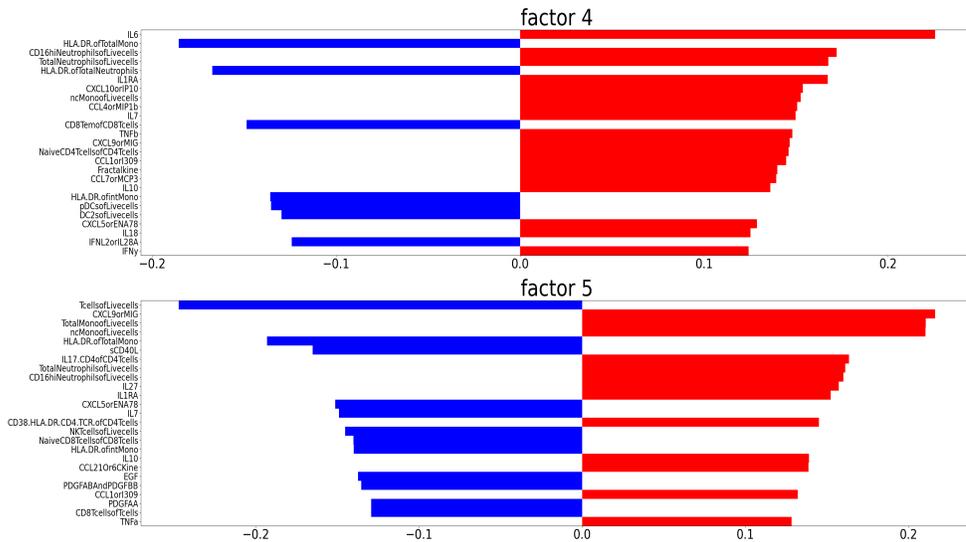
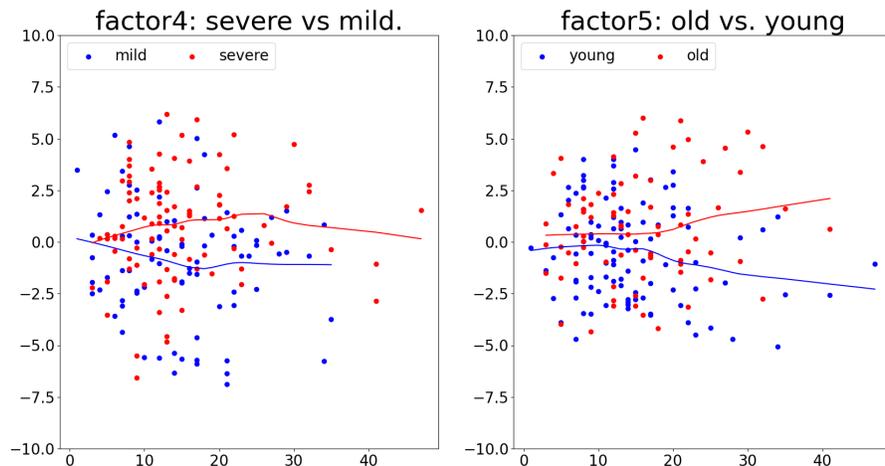


FIG 6. Weighted time trajectories for different patient groups. Left: Comparison of mild and severe patient groups, with features weighted by factor 4. Right: Comparison of young and old patient groups, with features weighted by factor 5. Horizontal time axes show the DFSO, and y-axes show the weighted feature values.



We mention two useful extensions of SPACO here. We can encourage sparsity in V by imposing a lasso penalty in the estimation of V . This sparsity may be helpful for interpretability when features are high-dimensional, but only some are relevant to a given factor. It is also interesting to extend SPACO to model multi-omics data. Such data may introduce two complications: First, different omics can be measured at different times. Second, the data type and data scale can be very different. For example, the number of features for RNA-seq

data (Marguerat and Bähler, 2010) is often of much higher dimension than the number of cell types derived from single-cell experiments (Trapnell, 2015), and DNA methylation data is binary instead of Gaussian (Bock, 2012). These reasons motivate the more careful design of a tailored model, rather than a naive approach of blindly pooling different omics data together.

APPENDIX

We provide proofs to Propositions, Lemmas, and Theorems in the main paper in Appendix A. In Appendix B, we describe the estimation details in procedures left out in the main paper. We also include qq-plots comparing p-values from cross-fitted and non-cross-fitted SPACO in Appendix C.

APPENDIX A: PROOFS

APPENDIX: PROOF OF PROPOSITIONS

A.1. Proof of Proposition 2.1.

PROOF. From eq. (5), the marginal log-likelihood for any subject i is

$$2L(\mathbf{X}_{I,\vec{i}}|\Theta) \propto -\mathbf{z}_i^\top \boldsymbol{\beta}^\top \Lambda_f \boldsymbol{\beta} \mathbf{z}_i + \boldsymbol{\mu}_i^\top \Sigma_i \boldsymbol{\mu}_i + \log |\Lambda_f^{\frac{1}{2}} \Sigma_i \Lambda_f^{\frac{1}{2}}|.$$

When we plug in the expression for the posterior covariance Σ_i and means $\boldsymbol{\mu}_i$ in to the above expression, we have

$$\Lambda_f^{\frac{1}{2}} \Sigma_i \Lambda_f^{\frac{1}{2}} = \mathbf{Id} + \Lambda_f^{\frac{1}{2}} (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} (\mathbf{V} \odot \boldsymbol{\Phi})_i \Lambda_f^{\frac{1}{2}},$$

and

$$\begin{aligned} & \boldsymbol{\mu}_i^\top \Sigma_i \boldsymbol{\mu}_i \\ &= \left(\Lambda_f \boldsymbol{\beta}^\top \mathbf{z}_i + (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} \mathbf{X}_{I,\vec{i}} \right)^\top \left(\Lambda_f + (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} (\mathbf{V} \odot \boldsymbol{\Phi})_i \right)^{-1} \left(\Lambda_f \boldsymbol{\beta}^\top \mathbf{z}_i + (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} \mathbf{X}_{I,\vec{i}} \right) \\ &= \mathbf{z}_i^\top \boldsymbol{\beta} \Lambda_f^{\frac{1}{2}} \left(\mathbf{Id} + \Lambda_f^{-\frac{1}{2}} (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} (\mathbf{V} \odot \boldsymbol{\Phi})_i \Lambda_f^{-\frac{1}{2}} \right)^{-1} \Lambda_f^{\frac{1}{2}} \boldsymbol{\beta}^\top \mathbf{z}_i \\ & \quad + \mathbf{X}_{I,\vec{i}}^\top \Lambda_i^{-1} (\mathbf{V} \odot \boldsymbol{\Phi})_i \Lambda_f^{-\frac{1}{2}} \left(\mathbf{Id} + \Lambda_f^{-\frac{1}{2}} (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} (\mathbf{V} \odot \boldsymbol{\Phi})_i \Lambda_f^{-\frac{1}{2}} \right)^{-1} \Lambda_f^{-\frac{1}{2}} (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} \mathbf{X}_{I,\vec{i}} \\ & \quad + 2 \mathbf{z}_i^\top \boldsymbol{\beta} \Lambda_f^{\frac{1}{2}} \left(\mathbf{Id} + \Lambda_f^{-\frac{1}{2}} (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} (\mathbf{V} \odot \boldsymbol{\Phi})_i \Lambda_f^{-\frac{1}{2}} \right)^{-1} \Lambda_f^{-\frac{1}{2}} (\mathbf{V} \odot \boldsymbol{\Phi})_i^\top \Lambda_i^{-1} \mathbf{X}_{I,\vec{i}}. \end{aligned}$$

Let C_1, C_2, C_3 be $K \times K$ diagonal matrix with k^{th} diagonal entry being c_1, c_2, c_3 respectively and other diagonal entries being 1. When $c_1 c_2 c_3 = 1$, we have,

$$\begin{aligned} & \Lambda_f^{-\frac{1}{2}} C_1 \left((\mathbf{V} C_3) \odot (\boldsymbol{\Phi} C_2) \right) = \Lambda_f^{-\frac{1}{2}} C_1 C_2 C_3 \left(\mathbf{V} \odot \boldsymbol{\Phi} \right) = \Lambda_f^{-\frac{1}{2}} \left(\mathbf{V} \odot \boldsymbol{\Phi} \right) \\ & \mathbf{z}_i^\top (\boldsymbol{\beta} C_1) (C_1^{-1} \Lambda_f^{\frac{1}{2}}) = \mathbf{z}_i^\top \boldsymbol{\beta} \Lambda_f^{\frac{1}{2}}. \end{aligned}$$

As a result, when $c_1 c_2 c_3 = 1$, we have $L(X_{I,\vec{i}}|\Theta_k(c_1, c_2, c_3), \Theta_{k^c}) = L(X_{I,\vec{i}}|\Theta_k(1, 1, 1), \Theta_{k^c})$ for all $i = 1, \dots, I$. Consequently, we have $L(X_I|\Theta_k(c_1, c_2, c_3), \Theta_{k^c}) = L(X_I|\Theta_k(1, 1, 1), \Theta_{k^c})$. \square

A.2. Proof of Proposition 4.2.

PROOF. It is immediate from Lemme 3.1 that SPACO solves for β_k considering

$$J(\beta_k) = \min_{\beta_k} \left\{ \sum_{i=1}^I \frac{1}{w_i} (\tilde{\mathbf{z}}_i^\top \beta_k - \tilde{y}_i)^2 + \lambda_{2k} |\beta_k|_1 \right\},$$

where $w_i = s_k^2 - (\Sigma_i)_{kk}$, $\tilde{\mathbf{z}}_i = (1 - \frac{(\Sigma_i)_{kk}}{s_k^2}) \mathbf{z}_i$, and $\tilde{y}_i = \left((\Sigma_i)_{k,:} (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \Lambda_{\tilde{i}} \mathbf{X}_{I,\tilde{i}} + \sum_{\ell \neq k} \frac{(\Sigma_i)_{kl}}{s_\ell^2} \mathbf{z}_i^\top \beta_\ell^* \right)$.

Recall that $\mathbf{X}_{I,\tilde{i}} = (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \mathbf{u}_i = (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top (\beta^\top \mathbf{z}_i + \zeta_i) + \epsilon_{I,\tilde{i}}$, and $\Sigma_i = \left(\Lambda_f + (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \Lambda_{\tilde{i}} (\mathbf{V} \odot \Phi)_{\tilde{i}} \right)^{-1}$.

Hence, let $\mathbf{e}_k = (\underbrace{0, \dots, 0}_{k-1}, 1, \underbrace{0, \dots, 0}_{K-k})^\top$, we have

$$\begin{aligned} (\Sigma_i)_{k,:} (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \Lambda_{\tilde{i}} \mathbf{X}_{I,\tilde{i}} &= (\mathbf{e}_k^\top - (\Sigma_i)_{k,:} \Lambda_f) \beta^{*T} \mathbf{z}_i + (\mathbf{e}_k^\top - (\Sigma_i)_{k,:} \Lambda_f) \zeta_i + (\Sigma_i)_{k,:} (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \Lambda_{\tilde{i}} \epsilon_{I,\tilde{i}} \\ &= (1 - \frac{(\Sigma_i)_{kk}}{s_k^2}) \mathbf{z}_i^\top \beta_k^* - \sum_{\ell \neq k} \frac{(\Sigma_i)_{kl}}{s_\ell^2} \mathbf{z}_i^\top \beta_\ell^* + \xi_i \end{aligned}$$

where $\xi_i = (\mathbf{e}_k^\top - \Sigma_{ik} \Lambda_f) \zeta_i + \Sigma_{ik} (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \Lambda_{\tilde{i}} \epsilon_{I,\tilde{i}}$, and

$$\begin{aligned} \mathbb{E}(\xi_i^2) &= (\mathbf{e}_k^\top - \Sigma_{ik} \Lambda_f) \Lambda_f^{-1} (\mathbf{e}_k - \Lambda_f (\Sigma_i)_{:,k}) + (\Sigma_i)_{k,:} (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top \Lambda_{\tilde{i}} (\mathbf{V} \odot \Phi)_{\tilde{i}}^\top (\Sigma_i)_{:,k} \\ &= s_k^2 - 2(\Sigma_i)_{kk} + (\Sigma_i)_{k,:} \Lambda_f (\Sigma_i)_{:,k} + (\Sigma_i)_{k,:} (\Sigma_i^{-1} - \Lambda_f) (\Sigma_i)_{:,k} = s_k^2 - (\Sigma_i)_{kk} = w_i \end{aligned}$$

Consequently, we have $\tilde{y}_i = \tilde{\mathbf{z}}_i^\top \beta_k + \zeta_i$, with $\zeta_i \sim \mathcal{N}(0, w_i)$. \square

APPENDIX: PROOFS OF LEMMAS AND THEOREMS

A.3. Proof of Lemma 3.1. Lemma A.1 will be used for the derivation of the proximal update and the proof of the objective being non-decreasing,

LEMMA A.1. *Let A be a semi-positive definite matrix. For the constrained problem*

$$\min_{x: \|x\|_2^2 = h_0} f(x) := \frac{1}{2} x^\top A x - a^\top x,$$

we define the proximal mapping at x_0 and with step size ρ as

$$\text{prox}_\rho(x_0) = \begin{cases} \frac{1}{2} \|x - x_0\|_2^2 + \rho(Ax_0 - a) & \text{if } \|x\|_2^2 = h_0 \\ \infty & \text{otherwise} \end{cases}$$

The solution to $\text{prox}_\rho(x_0)$ takes the form

$$x^* = \frac{\sqrt{h_0} \tilde{x}}{\|\tilde{x}\|_2}, \quad \tilde{x} = x_0 - \rho(Ax_0 - a).$$

Let γ be the largest eigenvalue of A . When $\|x_0\|_2^2 = h_0$ and $\rho \leq \frac{1}{\gamma}$, we have $f(x^) \leq f(x_0)$.*

The proof of Lemma A.1 is given in Appendix A.8.

PROOF. We now show the derivation of the updating formulas.

Update β . We update β_k iteratively for $k = 1, \dots, K$ using the marginal log likelihood fixing all other parameters. Given the model parameters, the marginal log likelihood can be expressed as

$$L(\mathbf{X}|\Theta) \propto -\frac{1}{2} \sum_i \left(\mathbf{X}_{I,\vec{i}}^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} + \mathbf{z}_i^\top \beta \Lambda_f \beta^\top \mathbf{z}_i - \mu_i^\top \Sigma_i^{-1} \mu_i \right) + \frac{1}{2} \left(\sum_i (\log |\Lambda_{\vec{i}}| + \log |\Sigma_i|) + I \log |\Lambda_f| \right)$$

where $\Sigma_i = \left(\Lambda_f + (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} (\mathbf{V} \odot \Phi)_{\vec{i}} \right)^{-1}$, $\mu_i = \Sigma_i \left(\Lambda_f \beta^\top \mathbf{z}_i + (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} \right)$ are the posterior covariance and mean of \mathbf{u}_i . As a consequence, maximizing the penalized marginal log likelihood with respect to β_k is equivalent to minimizing the following lasso penalized squared error loss:

$$J(\beta_k) = \frac{1}{w_i} \left(\left(\frac{s_k^2 - (\Sigma_i)_{kk}}{s_k^2} \right)^2 \beta_k^\top \mathbf{z}_i \mathbf{z}_i^\top \beta_k - 2 \left(\frac{s_k^2 - (\Sigma_i)_{kk}}{s_k^2} \right) \beta_k^\top \mathbf{z}_i \left((\Sigma_i)_{k\cdot} (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} + (\Sigma_i)_{kk} \Lambda_{f,k^c} \beta_{k^c}^\top \mathbf{z}_i \right) \right) + \lambda_{2k} |\beta_k|,$$

where $w_i = s_k^2 - (\Sigma_i)_{kk}$. Let $\tilde{\mathbf{z}}_i = \sqrt{\frac{1}{s_k^2} \frac{s_k^2 - (\Sigma_i)_{kk}}{s_k^2}} \mathbf{z}_i$, and

$$\tilde{y}_i = \sqrt{\frac{1}{s_k^2} \frac{s_k^2}{s_k^2 - (\Sigma_i)_{kk}}} \left((\Sigma_i)_{i,k} (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} + (\Sigma_i)_{kk} \Lambda_{f,k^c} \beta_{k^c}^\top \mathbf{z}_i \right).$$

Then minimization of $J(\beta_k)$ is equivalent to solving the following lasso penalized regression problem

$$\min_{\beta_k} \frac{1}{2} \|\tilde{\mathbf{y}} - \tilde{\mathbf{Z}} \beta_k\|_2^2 + \lambda_{2k} |\beta_k|,$$

where $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_I)$ and $\tilde{\mathbf{Z}}$ is the row stack of $\tilde{\mathbf{z}}_i$.

Update s^2 . Let $\mathbb{E}_{\Theta_0} L(\mathbf{X}, \mathbf{U}|\Theta)$ be the expected log likelihood under $P(\mathbf{U}|\Theta_0)$ and the part of it that relates to $\{s_k^2, k = 1, \dots, K\}$ is

$$\begin{aligned} 2\mathbb{E}_{\Theta_0} L(\mathbf{X}, \mathbf{U}|\Theta) &\propto \sum_{i=1}^I \left\{ \mathbb{E}_{\Theta_0} [-(\mathbf{u}_i - \beta^\top \mathbf{z}_i)^\top \Lambda_f (\mathbf{u}_i - \beta^\top \mathbf{z}_i)] + \log |\Lambda_f| \right\} \\ &\propto \sum_{k=1}^K \left\{ \frac{\sum_{i=1}^I ((\mu_{ik} - \mathbf{z}_i^\top \beta_k)^2 + (\Sigma_i)_{kk})}{s_k^2} - I \log s_k^2 \right\} \end{aligned}$$

Consequently, the solution for s_k^2 given other parameters is $s_k^2 = \frac{1}{I} \sum_{i=1}^I ((\mu_{ik} - \mathbf{z}_i^\top \beta_k)^2 + (\Sigma_i)_{kk})$.

Update σ^2 . The part of expected log likelihood related to \mathbf{V} and σ^2 can be expressed as

$$\begin{aligned} J(\mathbf{V}, \sigma^2) &= -\frac{1}{2} \sum_{j=1}^J \left\{ \frac{1}{\sigma_j^2} \mathbb{E}_{\Theta_0} (X_{J,\vec{j}} - (\Phi \odot \mathbf{U})_{\vec{j}} \mathbf{v}_j)^\top (X_{J,\vec{j}} - (\Phi \odot \mathbf{U})_{\vec{j}} \mathbf{v}_j) + |\vec{j}| \log \sigma_j^2 \right\} \\ &= -\frac{1}{2} \sum_{j=1}^J \left\{ \frac{1}{\sigma_j^2} \left[(X_{J,\vec{j}} - (\Phi \odot \mu)_{\vec{j}} \mathbf{v}_j)^\top (X_{J,\vec{j}} - (\Phi \odot \mu)_{\vec{j}} \mathbf{v}_j) + \mathbf{v}_j^\top \left(\sum_t \left(\sum_{O_{it}=1} \Sigma_i \right) \cdot (\phi_t \phi_t^\top) \right) \mathbf{v}_j \right] + |\vec{j}| \log \sigma_j^2 \right\} \end{aligned}$$

Hence, our update rule for σ_j^2

$$\sigma_j^2 = \frac{1}{|\vec{j}|} \left[(X_{J,\vec{j}} - (\Phi \odot \mu)_{\vec{j}} \mathbf{v}_j)^\top (X_{J,\vec{j}} - (\Phi \odot \mu)_{\vec{j}} \mathbf{v}_j) + \mathbf{v}_j^\top \left(\sum_t \left(\sum_{O_{it}=1} \Sigma_i \right) \cdot (\phi_t \phi_t^\top) \right) \mathbf{v}_j \right].$$

Update \mathbf{V} . Define $M^j = \frac{1}{\sigma_j^2} \left[(\Phi \odot \mu)_{\bar{j}} (\Phi \odot \mu)_{\bar{j}}^\top + (\sum_t (\sum_{O_{it}=1} \Sigma_i) \cdot (\phi_t \phi_t^\top)) \right]$, and $H^j = \frac{1}{\sigma_j^2} (\Phi \odot \mu)_{\bar{j}}^\top X_{J,\bar{j}}$. The expected negative log likelihood associated with \mathbf{V}_k can be written as

$$J(\mathbf{V}_k) = \frac{1}{2} \mathbf{V}_k^\top A \mathbf{V}_k - a^\top \mathbf{V}_k,$$

where $A = \text{diag}\{M_{kk}^1, \dots, M_{kk}^J\}$ and $a = H_k^{1:J} - \mathbf{V}_k^c M_{k^c,k}^{1:J}$. Our proximal mapping for \mathbf{V}_k at the current parameter value and step size ρ can be written as

$$\text{prox}_{\mathbf{V}_k}(\mathbf{y}; \mathbf{V}_k, \rho) = \begin{cases} \frac{1}{2} \|\mathbf{y} - \mathbf{V}_k\|_2^2 - \rho(a - A\mathbf{V}_k)^\top \mathbf{y} & \text{if } \|\mathbf{y}\|_2^2 = 1 \\ \infty & \text{otherwise,} \end{cases}$$

By Lemma A.1, the solution to the proximal mapping is

$$\mathbf{V}_k \leftarrow \mathbf{y}^* = \arg \min_{\mathbf{y}} \text{prox}_{\mathbf{V}_k}(\mathbf{y}; \mathbf{V}_k, \alpha) = \frac{\tilde{\mathbf{y}}}{\|\tilde{\mathbf{y}}\|_2}, \quad \tilde{\mathbf{y}}_j = \mathbf{V}_{jk} + \rho(a - A\mathbf{V}_k).$$

Update Φ . The part of expected log likelihood related to Φ can be expressed as

$$\begin{aligned} J(\Phi) &= \sum_{t=1}^T \mathbb{E}_{\Theta_0} \left[(X_{T,\bar{t}} - (\mathbf{V} \odot \mathbf{U})_{\bar{t}} \phi_t)^\top \Lambda_{\bar{t}} (X_{T,\bar{t}} - (\mathbf{V} \odot \mathbf{U})_{\bar{t}} \phi_t) \right] \\ &\propto \sum_{t=1}^T \left\{ \phi_t^\top \mathbb{E}_{\Theta_0} \left[(\mathbf{V} \odot \mathbf{U})_{\bar{t}}^\top \Lambda_{\bar{t}}^{-1} (\mathbf{V} \odot \mathbf{U})_{\bar{t}} \right] \phi_t - 2\phi_t^\top \mathbb{E}_{\Theta_0} \left[(\mathbf{V} \odot \mathbf{U})_{\bar{t}}^\top \Lambda_{\bar{t}} \right] X_{T,\bar{t}} \right\} \\ &= \sum_{t=1}^T \left\{ \phi_t^\top \left[(\mathbf{V} \odot \mu)_{\bar{t}}^\top \Lambda_{\bar{t}} (\mathbf{V} \odot \mu)_{\bar{t}} + \sum_j \frac{(\sum_{O_{it}=1} \Sigma_i) \cdot (\mathbf{v}_j \mathbf{v}_j^\top)}{\sigma_j^2} \right] \phi_t - 2\phi_t^\top \left[(\mathbf{V} \odot \mu)_{\bar{t}}^\top \Lambda_{\bar{t}} \right] X_{T,\bar{t}} \right\} \end{aligned}$$

Define

$$M^t = \left[(\mathbf{V} \odot \mu)_{\bar{t}}^\top \Lambda_{\bar{t}}^{-1} (\mathbf{V} \odot \mu)_{\bar{t}} + \sum_j \frac{(\sum_{O_{it}=1} \Sigma_i) \cdot (\mathbf{v}_j \mathbf{v}_j^\top)}{\sigma_j^2} \right], \quad H^t = (\mathbf{V} \odot \mu)_{\bar{t}}^\top \Lambda_{\bar{t}} X_{T,\bar{t}}.$$

Then, together with the smoothness penalty, we are considering the following constrained quadratic problem when updating Φ_k :

$$\min_{\mathbf{V}_k} \frac{1}{2} \mathbf{V}_k^\top (A + \lambda_{1k} \Omega) \mathbf{V}_k - a^\top \mathbf{V}_k,$$

where $A = \text{diag}\{M_{kk}^1, \dots, M_{kk}^T\}$ and $a \in \mathbb{R}^T$ is a stack of $H_k^t - \Phi_{k^c} M_{k^c,k}^t$ for $t = 1, \dots, T$. Same as in the derivation of updating \mathbf{V}_k , by Lemma A.1, we have

$$\Phi_k \leftarrow \mathbf{y}^* = \arg \min_{\mathbf{y}} \text{prox}_{\Phi_k}(\mathbf{y}; \Phi_k, \rho) = \frac{\tilde{\mathbf{y}}}{\|\tilde{\mathbf{y}}\|_2}, \quad \tilde{\mathbf{y}} = \Phi_k + \rho(a - A\mathbf{V}_k - \lambda_{1k} \Omega \mathbf{V}_k).$$

□

A.4. Proof of Theorem 3.2.

PROOF. The update from Θ^ℓ to $\Theta^{\ell+1}$ consists of two steps:

- $\Theta^\ell \Rightarrow (\Theta^\ell \setminus \{\beta^\ell\}) \cup \{\beta^{\ell+1}\}$.
- $(\Theta^\ell \setminus \{\beta^\ell\}) \cup \{\beta^{\ell+1}\} \Rightarrow \Theta^{\ell+1}$.

The first step leads to non-decreasing marginal log-likelihood by definition, while the second step is a EM procedure. If we can show that the penalized marginal log-likelihood is non-decreasing at each subroutine of the EM procedure, we prove Theorem 3.2.

For simplicity, θ be the parameters that is being updated in some subroutine and $\Theta_{\setminus\theta}$ parameters excluding θ . Let $\Theta = \Theta_{\setminus\theta} \cup \theta$, $\Theta' = \Theta_{\setminus\theta} \cup \theta'$. It is known that if a new θ' is no worse compared with θ using the EM objective, it is no worse than θ when it comes to the (regularized) marginal MLE (Dempster, Laird and Rubin, 1977). We include the short proof here for completeness.

Because the log of the posterior of U can be decomposed into the difference between the log of the log complete likelihood and the log marginal likelihood, $L(U|X, \Theta) = L(U, X|\Theta) - L(X|\Theta)$, we have (expectation with respect to posterior distribution of U with parameters Θ):

$$\mathbb{E}_{\Theta} L(U, X|\Theta') = \mathbb{E}_{\Theta} L(X|\Theta') + \mathbb{E}_{\Theta} L(U|X, \Theta') = L(X|\Theta') + \mathbb{E}_{\Theta} L(U|X, \Theta').$$

As a result, when

$$\mathbb{E}_{\Theta} L(X, U|\Theta') \geq \mathbb{E}_{\Theta} L(X, U|\Theta),$$

the following inequality holds,

$$\begin{aligned} L(X|\Theta') - L(X|\Theta) &= \{\mathbb{E}_{\Theta} L(X|\Theta') + \mathbb{E}_{\Theta} L(U|X, \Theta') - L(X|\Theta')\} + \{\mathbb{E}_{\Theta} L(U|X, \Theta') - \mathbb{E}_{\Theta} L(U|X, \Theta)\} \\ &\geq \mathbb{E}_{\Theta} \log \frac{P(U|X, \Theta)}{P(U|X, \Theta')} \geq 0. \end{aligned}$$

The last inequality is due to the fact that the mutual information $\mathbb{E}_{\Theta} \log \frac{P(U|X, \Theta)}{P(U|X, \Theta')}$ is nonnegative.

For our subroutines of updating s^2 and σ^2 , they are defined as the maximizers of $\mathbb{E}_{\Theta} L(X, U|\Theta_{-\theta}, \theta)$. For updating V_k and Φ_k , by Lemma A.1, we know that

- $\mathbb{E}_{\Theta} L(X, U|\Theta_{-\theta}, \theta)$ for $\theta = V_k$ if $\frac{1}{\rho_k(V)} \geq \max_j M_{kk}^j$.
- $\mathbb{E}_{\Theta} L(X, U|\Theta_{-\theta}, \theta)$ for $\theta = \Phi_k$ if $\frac{1}{\rho_k(\Phi)} \geq \text{eig}_{\max}(A_k^{\Phi} + \lambda_{1k}\Omega)$.

Hence, we have proved our statement. Hence, the log likelihood is non-decreasing or equivalently, the loss $J(\Theta)$ is non-increasing. \square

A.5. Proof of Lemma 4.1.

PROOF. If we can show that $X_I = U_{\perp} M_U^{\top}$, $X_T = \Phi_{\perp} M_T^{\top}$ and $X_T = V_{\perp}^{\top} M_V$. Then, there exists a rank- K core-array $G = \sum_{k=1}^K A_k \odot B_k \odot C_k$, such that

$$X = \sum_{k=1}^K U_k \odot \Phi_k \odot V_k,$$

where $U = U_{\perp} A$, $\Phi = \Phi_{\perp} B$ and $V = V_{\perp} C$. This statement can be checked easily: since U_{\perp} spans the row space of X_I , we can find a matrix A such that if we replace U^* with $U_{\perp} A$, we still have a PARAFAC decomposition of X . We can apply the same arguments to Φ_{\perp} , V_{\perp} , and hence prove the statement at the beginning.

The proposed V_{\perp} satisfies the requirement by construction. Hence, we need only to check that U_{\perp} and Φ_{\perp} spans the row space of X_I and X_T respectively.

- The projection of \mathbf{X}_J onto \mathbf{V}_\perp results in $\mathbf{H} = \mathbf{V}_\perp^\top \mathbf{X}_j = \mathbf{C}(\Phi^* \odot \mathbf{U}^*)^\top$, where $\mathbf{C} = \mathbf{V}_\perp^\top \mathbf{V}^*$. Hence, we have

$$\mathbf{W} = \sum_{k=1}^K \frac{1}{I} \mathbf{Y}(k)^\top \mathbf{Y}(k) = \Phi^* \left(\sum_{k=1}^K (\mathbf{U}^* \cdot \mathbf{C}_{k,\cdot})^\top (\mathbf{U}^* \cdot \mathbf{C}_{k,\cdot}) \right) \Phi^{*,T} = \Phi^* \mathbf{M} \Phi^{*\top}$$

where $\mathbf{M} = \mathbb{R}^{K \times K}$ with $M_{\ell k} = \langle \mathbf{C}_\ell, \mathbf{C}_k \rangle \langle \mathbf{U}_\ell^*, \mathbf{U}_k^* \rangle$. Notice that

$$\begin{aligned} \mathbf{X}_T \mathbf{X}_T^\top &= \Phi^* (\mathbf{V}^* \odot \mathbf{U}^*)^\top (\mathbf{V}^* \odot \mathbf{U}^*) \Phi^{*,T} \\ &= \Phi^* ((\mathbf{V}_\perp \mathbf{C}_1) \otimes \mathbf{U}_1^*, \dots, (\mathbf{V}_\perp \mathbf{C}_K) \otimes \mathbf{U}_K^*)^\top ((\mathbf{V}_\perp \mathbf{C}_1) \otimes \mathbf{U}_1^*, \dots, (\mathbf{V}_\perp \mathbf{C}_K) \otimes \mathbf{U}_K^*) \Phi^{*\top} \\ &= \Phi^* \left(\left((\mathbf{C}_\ell^\top \mathbf{V}_\perp^\top \mathbf{V}_\perp \mathbf{C}_k) \otimes (\mathbf{U}_\ell^\top \mathbf{U}_k) \right)_{\ell k} \right) \Phi^{*\top} \\ &= \Phi^* \left(\left(\mathbf{C}_\ell^\top \mathbf{C}_k \right) (\mathbf{U}_\ell^\top \mathbf{U}_k) \right)_{\ell k} \Phi^{*\top} = \Phi^* \mathbf{M} \Phi^{*\top} \end{aligned}$$

Hence, Φ_\perp is the top K left singular vectors of \mathbf{X}_T . We set $\mathbf{B} = \Phi_\perp^\top \Phi^*$.

- \mathbf{U}_\perp is estimated by regression \mathbf{X}_I on $\mathbf{V}_\perp \otimes \Phi$. Because both \mathbf{V}_\perp and Φ_\perp are orthonormal, $\mathbf{V}_\perp \otimes \Phi_\perp$ is also orthonormal:

$$(\mathbf{V}_\perp \otimes \Phi)^\top (\mathbf{V}_\perp \otimes \Phi) = (\mathbf{V}_\perp^\top \mathbf{V}_\perp) \otimes (\Phi^\top \Phi) = \mathbf{Id}_{K \times K} \otimes \mathbf{Id}_{K \times K} = \mathbf{Id}_{K^2 \times K^2}.$$

Hence, we have

$$\begin{aligned} \tilde{\mathbf{U}}(\mathbf{V}_\perp \otimes \Phi_\perp)^\top &= \frac{1}{1+\delta} \mathbf{X}_I (\mathbf{V}_\perp \otimes \Phi_\perp) (\mathbf{V}_\perp \otimes \Phi_\perp)^\top \\ &= \frac{1}{1+\delta} \mathbf{X}_I ((\mathbf{V}_\perp \mathbf{V}_\perp^\top) \otimes (\Phi_\perp \Phi_\perp^\top)) \\ &= \frac{1}{1+\delta} \mathbf{U}^* ((\mathbf{V}_\perp \mathbf{C}) \odot (\Phi_\perp \mathbf{B}))^\top (\mathbf{V}_\perp \otimes \Phi_\perp) (\mathbf{V}_\perp \otimes \Phi_\perp)^\top \\ &= \mathbf{U}^* \begin{pmatrix} (\mathbf{C}_1^\top \mathbf{V}_\perp^\top \mathbf{V}_\perp \mathbf{V}_\perp^\top) \otimes (\mathbf{B}_1^\top \Phi_\perp^\top \Phi_\perp \Phi_\perp^\top) & & \\ & \dots & \\ (\mathbf{C}_K^\top \mathbf{V}_\perp^\top \mathbf{V}_\perp \mathbf{V}_\perp^\top) \otimes (\mathbf{B}_K^\top \Phi_\perp^\top \Phi_\perp \Phi_\perp^\top) & & \end{pmatrix} (\mathbf{V}_\perp \otimes \Phi_\perp) \\ &= \frac{1}{1+\delta} \mathbf{X}_I \end{aligned}$$

The row space spanned by $\tilde{\mathbf{U}}$ is the same as the row space spanned by \mathbf{X}_I , thus, the space spanned by top K left singular vectors of $\tilde{\mathbf{U}}$ is the same by that of \mathbf{X}_I . As a result, \mathbf{U}_\perp also satisfies the requirement. In particular, we have

$$\mathbf{U}_\perp^\top \tilde{\mathbf{U}} = \frac{1}{1+\delta} \mathbf{U}_\perp^\top \mathbf{X}_I (\mathbf{V}_\perp \otimes \Phi_\perp) = \frac{1}{1+\delta} \mathbf{U}_\perp^\top \mathbf{U}_\perp \mathbf{G}_I (\mathbf{V}_\perp \otimes \Phi_\perp)^\top (\mathbf{V}_\perp \otimes \Phi_\perp) = \mathbf{G}_I$$

where \mathbf{G}_I is the unfolding of the core array \mathbf{G} in the subject dimension. Hence, we recover $\mathbf{A}, \mathbf{B}, \mathbf{C}$ applying a rank- K PARAFAC decomposition on the arranged three-dimensional core array from \mathbf{G}_I as described in Algorithm 2.

□

A.6. Proof of Lemma 4.3.

PROOF. Let $e_k = (\underbrace{0, \dots, 0}_{k-1}, 1, \underbrace{0, \dots, 0}_{K-k})^\top$. Plug-in the expression of $\mathbf{X}_{I,\vec{i}} = (\mathbf{V} \odot \Phi)_{\vec{i}} \mathbf{u}_i = (\mathbf{V} \odot \Phi)_{\vec{i}} ((\beta^*)^\top \mathbf{z}_i + \zeta_i) + \epsilon_{I,\vec{i}}$ into the expression of $\tilde{y}_i(\delta)$:

$$\begin{aligned} \tilde{y}_i(\delta) &= \left((\Sigma_i(\delta))_{k:} (\mathbf{V} \odot \Phi)^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} + \delta \sum_{\ell \neq k} \frac{(\Sigma_i(\delta))_{kl}}{s_\ell^2} \mathbf{z}_i^\top \beta_\ell \right) \\ &= (\mathbf{e}_k^\top - \delta (\Sigma_i(\delta))_{k:} \Lambda_f) (\beta^*)^\top \mathbf{z}_i + \delta \sum_{\ell \neq k} \frac{(\Sigma_i(\delta))_{kl}}{s_\ell^2} \mathbf{z}_i^\top \beta_\ell + (\mathbf{e}_k^\top - \delta (\Sigma_i(\delta))_{k:} \Lambda_f) \zeta_i \\ &\quad + (\Sigma_i(\delta))_{k:} (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} \epsilon_{I,\vec{i}} \\ &= (1 - \delta \frac{(\Sigma_i(\delta))_{kk}}{s_k^2}) \mathbf{z}_i^\top \beta_k^* - \sum_{\ell \neq k} \frac{(\Sigma_i(\delta))_{kl}}{s_\ell^2} \mathbf{z}_i^\top (\beta_\ell^* - \beta_\ell) + \xi_i \end{aligned}$$

where $\xi_i = (\mathbf{e}_k^\top - \delta (\Sigma_i(\delta))_{k:} \Lambda_f) \zeta_i + (\Sigma_i(\delta))_{k:} (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} \epsilon_{I,\vec{i}}$ and

$$\begin{aligned} \mathbb{E}(\xi_i^2) &= (\mathbf{e}_k^\top - \delta (\Sigma_i(\delta))_{k:} \Lambda_f) \Lambda_f^{-1} (\mathbf{e}_k - \delta \Lambda_f (\Sigma_i(\delta))_{:k}) + (\Sigma_i(\delta))_{k:} (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} (\mathbf{V} \odot \Phi)_{\vec{i}}^\top (\Sigma_i(\delta))_{:k} \\ &= s_k^2 - 2\delta (\Sigma_i(\delta))_{kk} + \delta^2 (\Sigma_i(\delta))_{k:} \Lambda_f (\Sigma_i(\delta))_{:k} (\delta) + (\Sigma_i(\delta))_{k:} (\Sigma_i^{-1} - \delta \Lambda_f) (\Sigma_i(\delta))_{:k} \\ &= s_k^2 + (1 - 2\delta) (\Sigma_i(\delta))_{kk} (\delta) + (\delta^2 - \delta) (\Sigma_i(\delta))_{k:} \Lambda_f (\Sigma_i(\delta))_{:k} = w_i(\delta) \end{aligned}$$

□

A.7. Proof of Lemma 4.4.

PROOF. Notice that

$$\tilde{y}_i(0) = \Sigma_{ik:}(0) (\mathbf{V} \odot \Phi)_{\vec{i}}^\top \Lambda_{\vec{i}} \mathbf{X}_{I,\vec{i}} = \mathbf{z}_i^\top \beta_k + \xi_i = \mu_i + \xi_i, \quad \xi_i \sim \mathcal{N}(0, w_i(0)).$$

The conclusions then follow from the following arguments:

1. Under the marginal null $H_{0k}^{marginal}$, we have $Z_j \perp\!\!\!\perp \mu_k$ and $Z_j \perp\!\!\!\perp \xi$, thus, $T_{marginal} | \tilde{\mathbf{y}}(0) \stackrel{d}{=} T_{marginal}^* | \tilde{\mathbf{y}}(0)$.
2. Under the conditional null $H_{0k}^{partial}$, we have $Z_j \perp\!\!\!\perp \mu_k | Z_{-j}$, $Z_j \perp\!\!\!\perp \xi$, and $\hat{\beta}_k$ is a fixed quantity given $\tilde{\mathbf{y}}(0)$ and \mathbf{Z}_{j^c} , and thus $T_{partial} | (\tilde{\mathbf{y}}(0), \mathbf{Z}_{j^c}) \stackrel{d}{=} T_{partial}^* | (\tilde{\mathbf{y}}(0), \mathbf{Z}_{j^c})$.

□

A.8. Proof of Lemma A.1.

PROOF. Suppose that the optimal solution to the proximal problem takes the form $x_* = \alpha \frac{\sqrt{h_0} \tilde{x}}{\|\tilde{x}\|_2} + \sqrt{(1 - \alpha^2) h_0} z$ for some $\alpha \in [0, 1]$ and z satisfying $z^\top \tilde{x} = 0$ and $\|z\|_2 = 1$. Then, let $C = \frac{1}{2} \|\tilde{x} - x_0\|_2^2 + \rho (Ax_0 - a)^\top \tilde{x}$, we have,

$$\begin{aligned} &\frac{1}{2} \|x_* - x_0\|_2^2 + \rho (Ax_0 - a)^\top x_* \\ &= C + \frac{1}{2} \|x_* - \tilde{x}\|_2^2 \\ &= C + (\tilde{x} - \alpha \frac{\sqrt{h_0}}{\|\tilde{x}\|_2} \tilde{x} - \sqrt{1 - \alpha^2} h_0 z)^2 \\ &= C + \|\tilde{x}\|_2^2 + \alpha^2 h_0 - 2\alpha \sqrt{h_0} \|\tilde{x}\|_2 + (1 - \alpha^2) h_0^2 = C + \|\tilde{x}\|_2^2 + h_0^2 - 2\alpha \sqrt{h_0} \|\tilde{x}\|_2 \end{aligned}$$

The smallest objective is reached at $\alpha = 1$, $x_* = \frac{\sqrt{h_0 \bar{x}}}{\|\bar{x}\|_2}$. For any x , let $f(x) = \frac{1}{2}x^\top Ax + b^\top x$, we have

$$f(x) \leq f(x_0) + f'(x_0)^\top (x - x_0) + \frac{1}{2} \text{eig}_{\max}(A) \|x - x_0\|_2^2$$

Since x_* is the optimal solution to the problem $\text{prox}_\rho(x_0)$, we have

$$f'(x_0)^\top (x_* - x_0) + \frac{1}{2\rho} \|x_* - x_0\|_2^2 \leq 0.$$

Combine them together:

$$f(x_*) - f(x_0) \leq f'(x_0)^\top (x_* - x_0) + \frac{1}{2\rho} \|x_* - x_0\|_2^2 + \frac{1}{2} \left(\frac{1}{\rho} - \text{eig}_{\max}(A) \right) \|x_* - x_0\|_2^2 \leq \frac{1}{2} \left(\frac{1}{\rho} - \lambda_{\max} \right) \|x_* - x_0\|_2^2.$$

As a result, if the step size $\rho \leq \frac{1}{\text{eig}_{\max}(A)}$, we are guaranteed to have $f(x_*) - f(x_0) \leq 0$. \square

APPENDIX B: ESTIMATIONS

B.1. Functional PCA for initializations. In Yao, Müller and Wang (2005), the authors suggest a functional PCA analysis by performing eigenvalue decomposition of the smoothed product matrix fitted with a local linear surface smoother. Here, we apply the suggested estimation approach on the total product matrix. Let $\hat{\mathbf{W}}_{i,s,t} = \sum_k \mathbf{Y}_{is}(k) \mathbf{Y}_{it}(k)$ be the empirical estimate of the total product matrix for subject i .

- To fit a local linear surface smoother for the off-diagonal element of \mathbf{W}_{s_0, t_0} , we consider the following problem:

$$\min \sum_i \sum_{O_{it} O_{is}=1, s \neq t} \kappa \left(\frac{s - s_0}{h_G}, \frac{t - t_0}{h_G} \right) (\hat{\mathbf{W}}_{i,s,t} - g((s_0, t_0), (s, t), \beta))^2,$$

with $g((s_0, t_0), (s, t), \beta) = \beta_0 + \beta_1(s - s_0) + \beta_2(t - t_0)$, and $\kappa: \mathbb{R}^2 \mapsto \mathbb{R}$ is a standard two-dimensional Gaussian kernel.

- For the diagonal element, we estimate it by local linear regression: for each t_0 :

$$\min \sum_i \sum_{O_{it}=1} \kappa_1 \left(\frac{t - t_0}{h_G} \right) (\hat{\mathbf{W}}_{i,t,t} - g(t_0, t, \beta))^2.$$

where $g(t_0, t, \beta) = \beta_0 + \beta_1(t - t_0)$.

By default, we let $h_G = \frac{T}{\sqrt{1 + \sum_{s \neq t} \mathbb{1}\{\sum_i O_{is} O_{it} > 0\}}}$.

B.2. Parameter tuning. In this section, we provide more details on the leave-one-time-out cross-validation for tuning λ_{1k} , $\forall k = 1, \dots, K$. The expected penalized deviance loss can be written as (keeping only terms relevant to Φ):

$$L = \sum_{t=1}^T \left\{ \phi_t^\top \left[(\mathbf{V} \odot \mu)_{\vec{t}}^\top \Lambda_{\vec{t}} (\mathbf{V} \odot \mu)_{\vec{t}} + \sum_j \frac{(\sum_{O_{it}=1} \Sigma_i) \cdot (\mathbf{v}_j \mathbf{v}_j^\top)}{\sigma_j^2} \right] \phi_t - 2\phi_t^\top \left[(\mathbf{V} \odot \mu)_{\vec{t}}^\top \Lambda_{\vec{t}} \right] X_{T, \vec{t}} \right\} + \sum_k \lambda_{1k} \Phi_k^\top \Omega \Phi_k.$$

For a given k , we define the diagonal matrix $\mathbf{A} \in \mathbb{R}^{T \times T}$ and the vector $a \in \mathbb{R}^T$ as

$$\mathbf{A}_{tt} = \left[(\mathbf{V} \odot \mu)_{\bar{t}}^\top \Lambda_{\bar{t}} (\mathbf{V} \odot \mu)_{\bar{t}} + \sum_j \frac{(\sum_{O_{it}=1} \Sigma_i) \cdot (\mathbf{v}_j \mathbf{v}_j^\top)}{\sigma_j^2} \right]_{kk}$$

$$a_t = (\mathbf{V}_k \otimes \mu_k)_{\bar{t}}^\top \Lambda_{\bar{t}} X_{T,\bar{t}} - \left\langle \left[(\mathbf{V} \odot \mu)_{\bar{t}}^\top \Lambda_{\bar{t}} (\mathbf{V} \odot \mu)_{\bar{t}} + \sum_j \frac{(\sum_{O_{it}=1} \Sigma_i) \cdot (\mathbf{v}_j \mathbf{v}_j^\top)}{\sigma_j^2} \right]_{k,k^c}, \phi_{t,k^c} \right\rangle.$$

When we leave out a specific time point t_0 , we optimize for Φ_k minimizing the following leave-one-out constrained loss,

$$\min_{\Phi_k} J(t_0, k) = \Phi_k^\top (\mathbf{A}(t_0) + \lambda_{1k} \Omega) \Phi_k - 2a(t_0)^\top \Phi_k, \quad \|\Phi_k\|_2^2 = T.$$

We set $\mathbf{A}(t_0)$ as \mathbf{A} with the (t_0, t_0) -entry zeroed out, and $a(t_0)$ as a with $a(t_0)$ zeroed out. The leave-one-time cross-validation error is calculated based on the expected deviance loss (unpenalized) at the leave-out time point t_0 :

$$J_{loocv}(t_0, k) = \mathbf{A}_{t_0, t_0} \phi_{t_0 k}^2 - 2a_{t_0} \phi_{t_0 k}.$$

The overall leave-one-time-out cross validation is calculated as $J_{loocv}(k) = \sum_{t=1}^T J_{loocv}(t, k)$.

At each step of updating Φ_k , we perform this this leave-one-column-out cross validation for every candidate penalty values of λ_{1k} , and select the penalty that achieves a small leave-one-column out error $L_{loocv}(k)$.

B.3. Generating Z_j^* . In the simulations and real data examples, we encounter two types of Z_j : Gaussian and binary. We model the conditional distribution of Z_j given Z_{j^c} with a (penalized) GLM. For Gaussian data, we consider a model where

$$Z_j = Z_{j^c} \theta + \epsilon_j, \quad \epsilon_j \sim \mathcal{N}(0, \sigma^2).$$

We estimate θ and σ^2 empirically from data. When q , the dimension of Z , is large, we apply a lasso penalty on β with penalty level selected with cross-validation. Let $\hat{\theta}$ and $\hat{\sigma}^2$ be our estimates of the distribution parameters. We then generate new \mathbf{z}_i^* for subject i from the estimated distribution $\mathbf{z}_i^* = \mathbf{z}_{i,j^c}^\top \hat{\theta} + \epsilon_{ij}^*$, with ϵ_{ij}^* independently generated from $\mathcal{N}(0, \hat{\sigma}^2)$. For binary Z_j , we consider the model

$$\log \frac{P(Z_j = 1)}{1 - P(Z_j = 1)} = Z_{j^c} \theta.$$

Again, we estimate θ empirically, with cross-validated lasso penalty for large q . We then generate z_{ij}^* independently from

$$P(Z_j = 1 | \mathbf{z}_{j^c}) = \frac{1}{1 + \exp(-\mathbf{z}_{i,j^c}^\top \hat{\theta})}.$$

To generate Z_j^* from the marginal distribution of Z_j , instead of estimating this distribution, we let \mathbf{Z}_j^* be a random permutation of \mathbf{Z}_j .

APPENDIX C: EMPIRICAL P-VALUES

In this section, we show qq-plots of the negative $\log_{10}(\text{pvalue})$ and compare the p-value distributions for SPACO with and without cross-fitting across 16 simulation settings discussed in our Section 5. The calibration for cross-fitted p-value is empirically good, whereas

the p-values derived without cross-fitting show a large deviation from uniform in quite a few settings. Like in the simulation setting, (SNR1, SNR2) are the signal to noise ratio in \mathbf{X} and \mathbf{Z} , J is the feature dimension of \mathbf{X} and "rate" is the observed rate in the time dimension. In all figures below, the horizontal axis shows the theoretical negative $\log_{10}(\text{p.value})$ and the vertical axis shows the observed negative $\log_{10}(\text{p.value})$.

FIG 7. *p*-values for partial associations using cross-fitting.

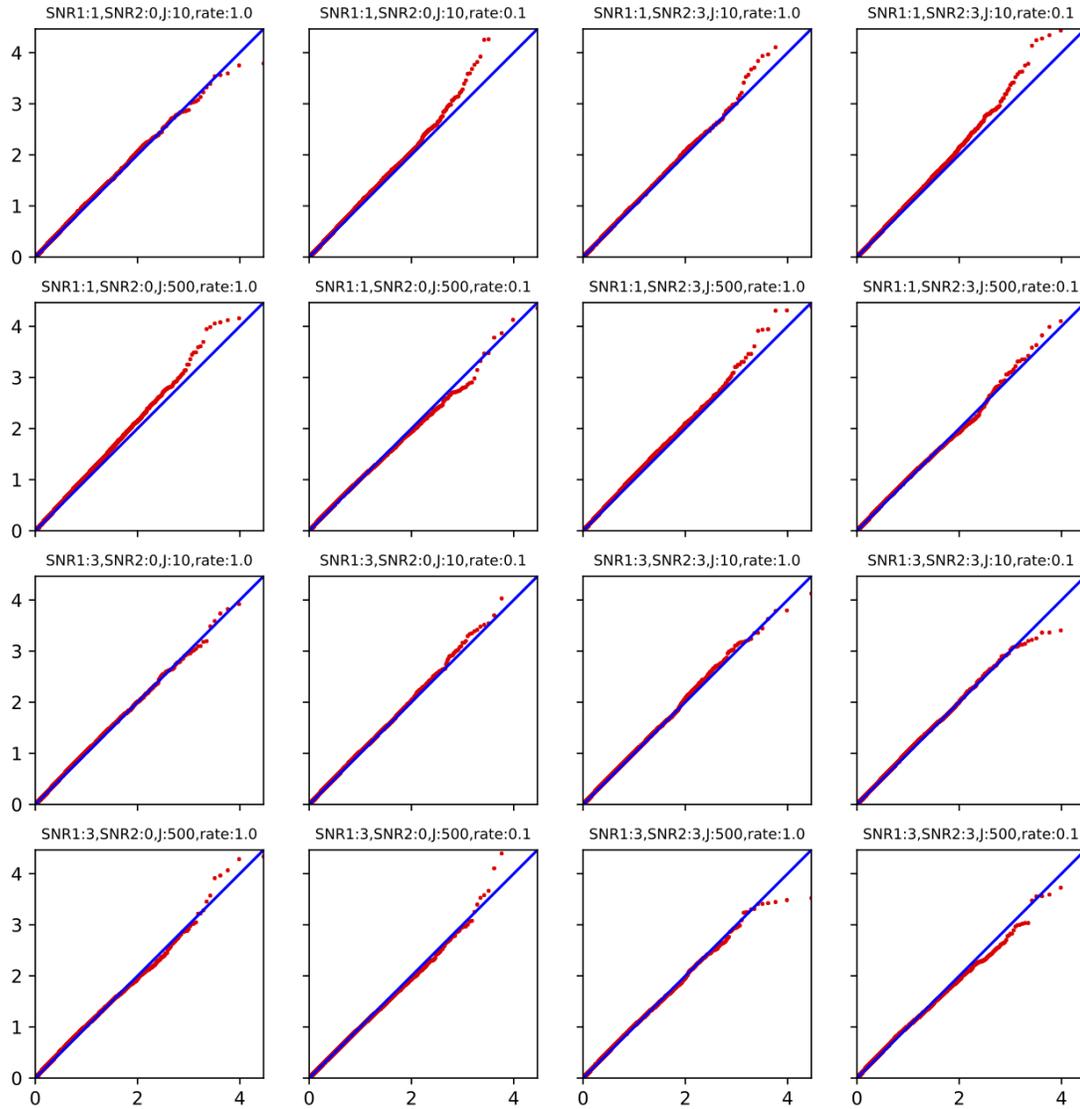


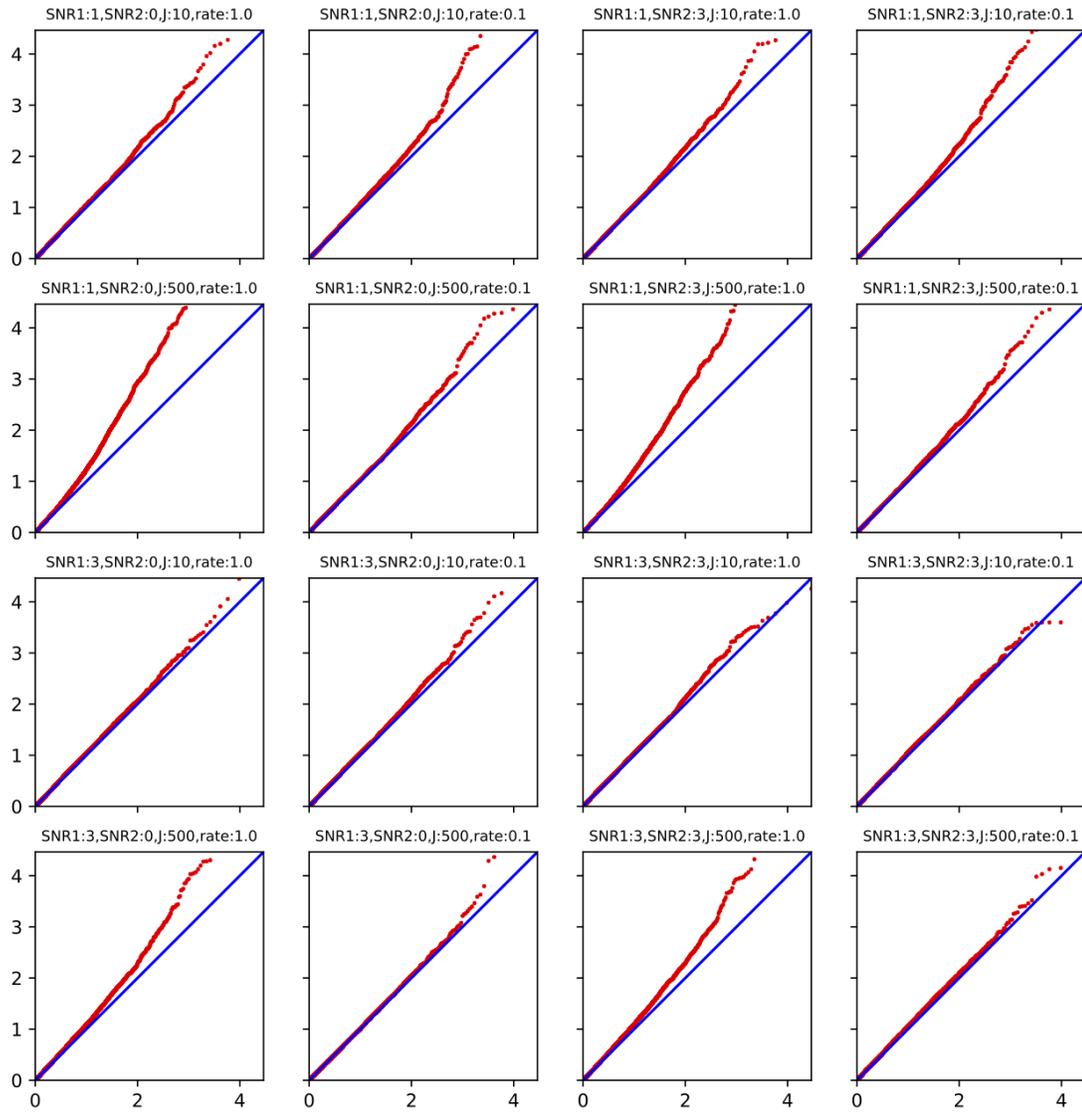
FIG 8. *p*-values for partial associations without cross-fitting.

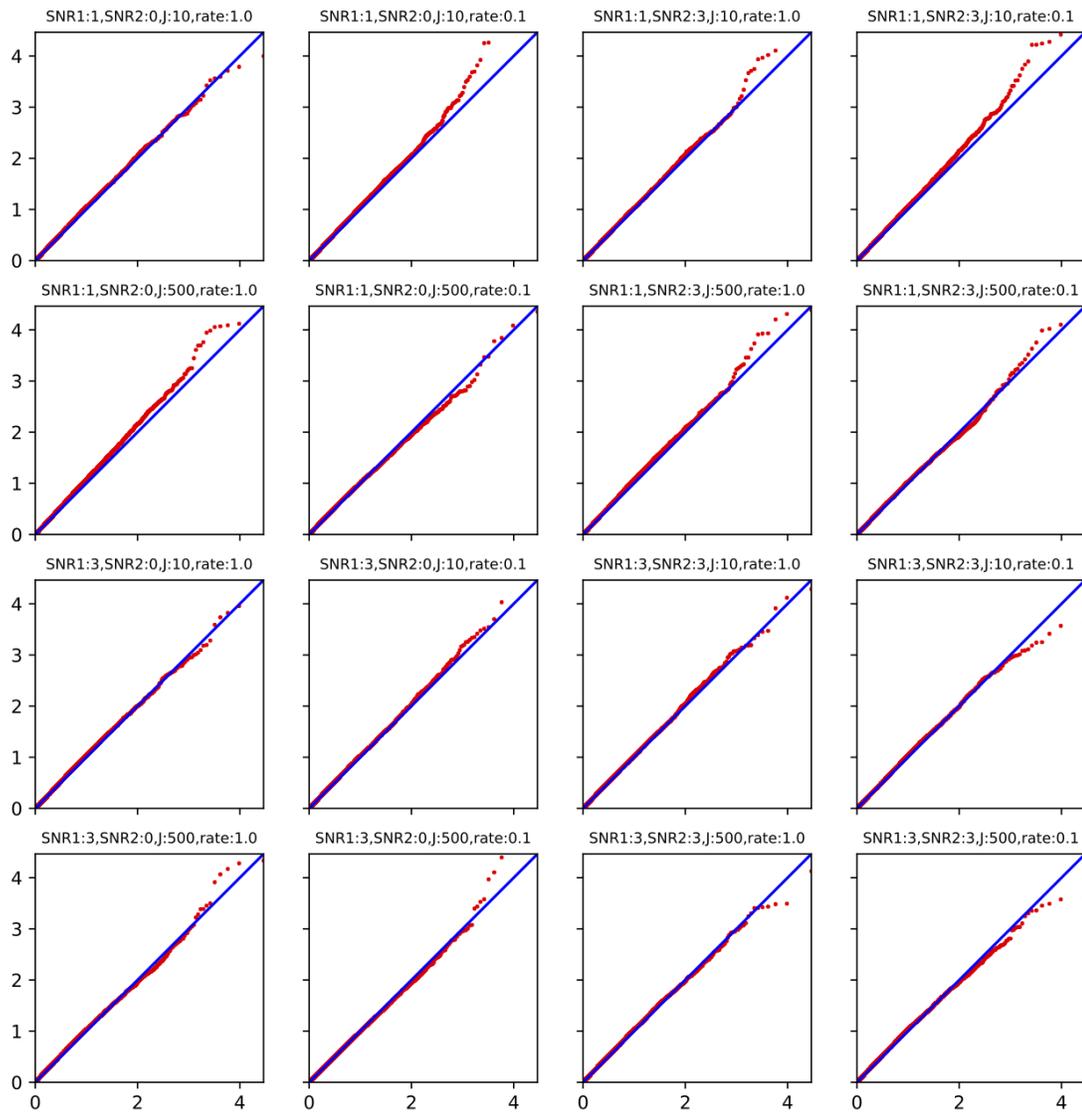
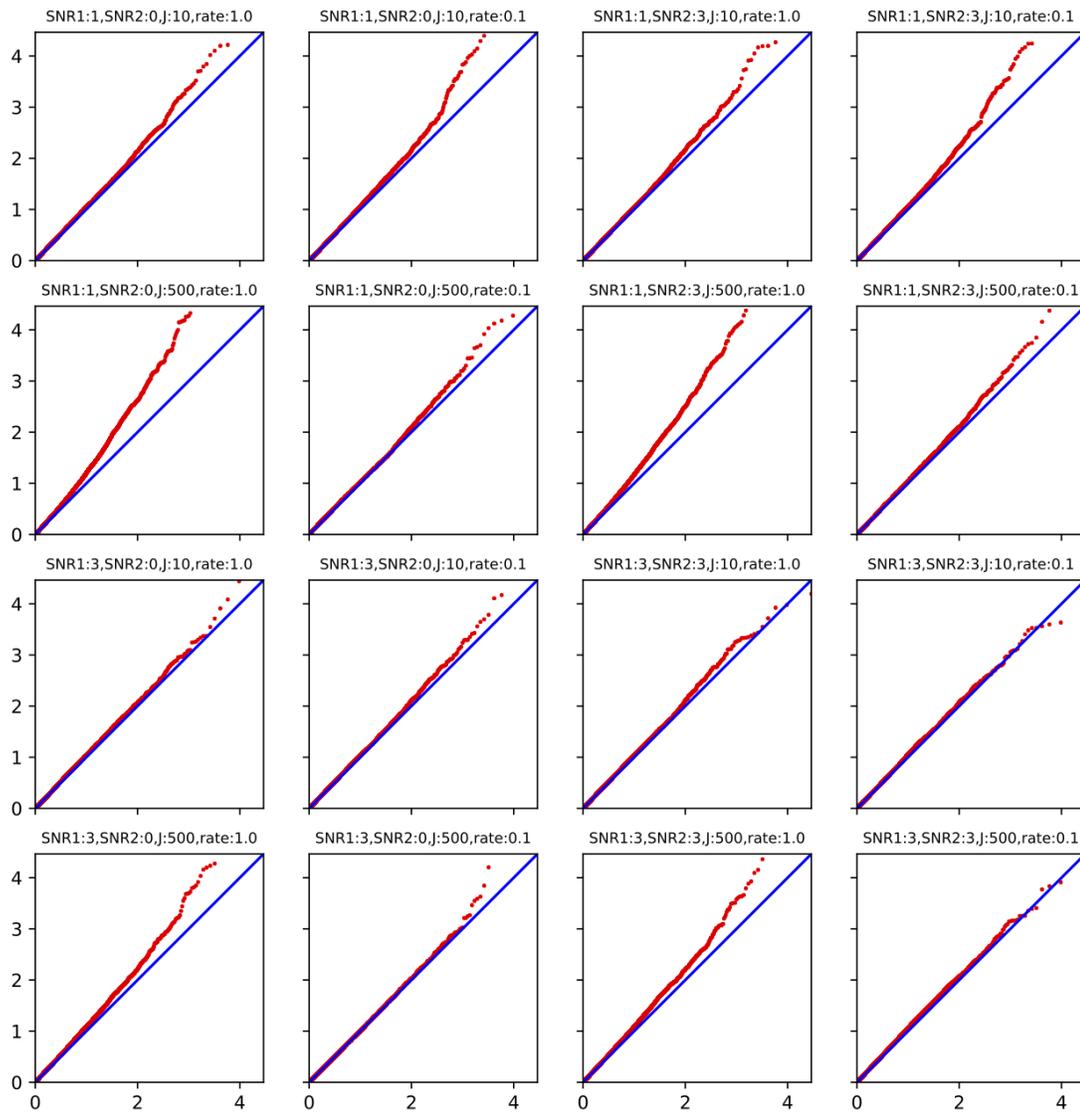
FIG 9. *p*-values for marginal associations using cross-fitting.

FIG 10. *p*-values for marginal associations without cross-fitting.

SUPPLEMENTARY MATERIAL

Code and data.

Code used for simulations and real data analysis, and the associated real data sets are available at https://github.com/LeyingGuan/spaco_py.git

REFERENCES

- ACAR, E. and YENER, B. (2008). Unsupervised multiway data analysis: A literature survey. *IEEE transactions on knowledge and data engineering* **21** 6–20.
- ANANDKUMAR, A., GE, R., HSU, D., KAKADE, S. M. and TELGARSKY, M. (2012). Tensor decompositions for learning latent variable models. *arXiv preprint arXiv:1210.7559*.
- ARGELAGUET, R., VELTEN, B., ARNOL, D., DIETRICH, S., ZENZ, T., MARIONI, J. C., BUETTNER, F., HUBER, W. and STEGLE, O. (2018). Multi-Omics Factor Analysis—a framework for unsupervised integration of multi-omics data sets. *Molecular systems biology* **14** e8124.
- BAIK, J. and SILVERSTEIN, J. W. (2006). Eigenvalues of large sample covariance matrices of spiked population models. *Journal of multivariate analysis* **97** 1382–1408.
- BESSE, P. and RAMSAY, J. O. (1986). Principal components analysis of sampled functions. *Psychometrika* **51** 285–311.
- BOCK, C. (2012). Analysing and interpreting DNA methylation data. *Nature Reviews Genetics* **13** 705–719.
- BRO, R. and ANDERSSON, C. A. (1998). Improving the speed of multiway algorithms: Part II: Compression. *Chemometrics and intelligent laboratory systems* **42** 105–113.
- CANDÈS, E., FAN, Y., JANSON, L. and LV, J. (2018). Panning for gold: ‘model-X’ knockoffs for high dimensional controlled variable selection Series B Statistical methodology.
- CARROLL, J. D., PRUZANSKY, S. and KRUSKAL, J. B. (1980). CANDELINC: A general approach to multidimensional analysis of many-way arrays with linear constraints on parameters. *Psychometrika* **45** 3–24.
- DE LATHAUWER, L., DE MOOR, B. and VANDEWALLE, J. (2000). A multilinear singular value decomposition. *SIAM journal on Matrix Analysis and Applications* **21** 1253–1278.
- DEMPSTER, A. P., LAIRD, N. M. and RUBIN, D. B. (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* **39** 1–22.
- DUCHENNE, O., BACH, F., KWEON, I.-S. and PONCE, J. (2011). A tensor-based algorithm for high-order graph matching. *IEEE transactions on pattern analysis and machine intelligence* **33** 2383–2395.
- FAN, J., FAN, Y. and LV, J. (2008). High dimensional covariance matrix estimation using a factor model. *Journal of Econometrics* **147** 186–197.
- FAN, J., LIAO, Y. and MINCHEVA, M. (2011). High dimensional covariance matrix estimation in approximate factor models. *Annals of statistics* **39** 3320.
- HARSHMAN, R. A. and LUNDY, M. E. (1994). PARAFAC: Parallel factor analysis. *Computational Statistics & Data Analysis* **18** 39–72.
- HINRICH, J. L. and MØRUP, M. (2019). Probabilistic tensor train decomposition. In *2019 27th European Signal Processing Conference (EUSIPCO)* 1–5. IEEE.
- HUANG, J. Z., SHEN, H., BUJA, A. et al. (2008). Functional principal components analysis via penalized rank one approximation. *Electronic Journal of Statistics* **2** 678–695.
- IMAIZUMI, M. and HAYASHI, K. (2017). Tensor decomposition with smoothness. In *International Conference on Machine Learning* 1597–1606. PMLR.
- JOHAN, H. (1990). Tensor Rank is NP-complete. *Journal of Algorithms* **4** 644–654.
- KATSEVICH, E. and ROEDER, K. (2020). Conditional resampling improves sensitivity and specificity of single cell CRISPR regulatory screens. *bioRxiv*.
- LAM, C., YAO, Q. and BATHIA, N. (2011). Estimation of latent factors for high-dimensional time series. *Biometrika* **98** 901–918.
- LI, G., SHEN, H. and HUANG, J. Z. (2016). Supervised sparse and functional principal component analysis. *Journal of Computational and Graphical Statistics* **25** 859–878.
- LUCAS, C., WONG, P., KLEIN, J., CASTRO, T. B., SILVA, J., SUNDARAM, M., ELLINGSON, M. K., MAO, T., OH, J. E., ISRAELOW, B. et al. (2020). Longitudinal analyses reveal immunological misfiring in severe COVID-19. *Nature* **584** 463–469.
- MARGUERAT, S. and BÄHLER, J. (2010). RNA-seq: from technology to biology. *Cellular and molecular life sciences* **67** 569–579.
- MESTRE, X. (2008). On the asymptotic behavior of the sample estimates of eigenvalues and eigenvectors of covariance matrices. *IEEE Transactions on Signal Processing* **56** 5353–5368.
- MNIH, A. and SALAKHUTDINOV, R. R. (2007). Probabilistic matrix factorization. *Advances in neural information processing systems* **20** 1257–1264.
- NION, D. and SIDIROPOULOS, N. D. (2010). Tensor algebra and multidimensional harmonic retrieval in signal processing for MIMO radar. *IEEE Transactions on Signal Processing* **58** 5693–5705.
- NION, D., MOKIOS, K. N., SIDIROPOULOS, N. D. and POTAMIANOS, A. (2009). Batch and adaptive PARAFAC-based blind separation of convolutive speech mixtures. *IEEE Transactions on Audio, Speech, and Language Processing* **18** 1193–1207.

- PAUL, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica* 1617–1642.
- PHAN, A.-H., TICHAVSKÝ, P. and CICHOCKI, A. (2013). CANDECOMP/PARAFAC decomposition of high-order tensors through tensor reshaping. *IEEE transactions on signal processing* **61** 4847–4860.
- RENDEIRO, A. F., CASANO, J., VORKAS, C. K., SINGH, H., MORALES, A., DESIMONE, R. A., ELLSWORTH, G. B., SOAVE, R., KAPADIA, S. N., SAITO, K. et al. (2020). Longitudinal immune profiling of mild and severe COVID-19 reveals innate and adaptive immune dysfunction and provides an early prediction tool for clinical progression. *medRxiv*.
- SIDIROPOULOS, N. D., DE LATHAUWER, L., FU, X., HUANG, K., PAPALEXAKIS, E. E. and FALOUTSOS, C. (2017). Tensor decomposition for signal processing and machine learning. *IEEE Transactions on Signal Processing* **65** 3551–3582.
- SORKINE, O., COHEN-OR, D., LIPMAN, Y., ALEXA, M., RÖSSL, C. and SEIDEL, H.-P. (2004). Laplacian surface editing. In *Proceedings of the 2004 Eurographics/ACM SIGGRAPH symposium on Geometry processing* 175–184.
- TAGUCHI, Y. (2019). Multiomics data analysis using tensor decomposition based unsupervised feature extraction. In *International Conference on Intelligent Computing* 565–574. Springer.
- TIBSHIRANI, R. (2011). Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73** 273–282.
- TIPPING, M. E. and BISHOP, C. M. (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **61** 611–622.
- TRAPNELL, C. (2015). Defining cell types and states with single-cell genomics. *Genome research* **25** 1491–1498.
- TUCKER, L. R. (1966). Some mathematical notes on three-mode factor analysis. *Psychometrika* **31** 279–311.
- WANG, D., LIU, X. and CHEN, R. (2019). Factor models for matrix-valued high-dimensional time series. *Journal of econometrics* **208** 231–248.
- WANG, D., ZHENG, Y., LIAN, H. and LI, G. (2021). High-dimensional vector autoregressive time series modeling via tensor decomposition. *Journal of the American Statistical Association* 1–19.
- XU, H., ZHANG, S., YI, X., PLEWCZYNSKI, D. and LI, M. J. (2020). Exploring 3D chromatin contacts in gene regulation: The evolution of approaches for the identification of functional enhancer-promoter interaction. *Computational and structural biotechnology journal* **18** 558–570.
- YAO, F., MÜLLER, H.-G. and WANG, J.-L. (2005). Functional data analysis for sparse longitudinal data. *Journal of the American statistical association* **100** 577–590.
- YOKOTA, T., ZHAO, Q. and CICHOCKI, A. (2016). Smooth PARAFAC decomposition for tensor completion. *IEEE Transactions on Signal Processing* **64** 5423–5436.