

Bidimensional linked matrix factorization for pan-omics pan-cancer analysis

Eric F. Lock¹, Jun Young Park¹, and Katherine A. Hoadley²

¹Division of Biostatistics, University of Minnesota

²Department of Genetics, University of North Carolina

Abstract

Several modern applications require the integration of multiple large data matrices that have shared rows and/or columns. For example, cancer studies that integrate multiple omics platforms across multiple types of cancer, *pan-omics pan-cancer analysis*, have extended our knowledge of molecular heterogeneity beyond what was observed in single tumor and single platform studies. However, these studies have been limited by available statistical methodology. We propose a flexible approach to the simultaneous factorization and decomposition of variation across such *bidimensionally linked* matrices, BIDIFAC+. This decomposes variation into a series of low-rank components that may be shared across any number of row sets (e.g., omics platforms) or column sets (e.g., cancer types). This builds on a growing literature for the factorization and decomposition of linked matrices, which has primarily focused on multiple matrices that are linked in one dimension (rows or columns) only. Our objective function extends nuclear norm penalization, is motivated by random matrix theory, gives an identifiable decomposition under relatively mild conditions, and can be shown to give the mode of a Bayesian posterior distribution. We apply BIDIFAC+ to pan-omics pan-cancer data from TCGA, identifying shared and specific modes of variability across 4 different omics platforms and 29 different cancer types.

1 Introduction

Data collection and curation for the Cancer Genome Atlas (TCGA) program completed in 2018, providing a unique and valuable public resource for comprehensive studies of molecular profiles across several types of cancer (Hutter and Zenklusen, 2018). The database includes information from several molecular platforms for over 10,000 tumor samples from individuals representing 33 types of cancer. The molecular platforms capture signal at different 'omics levels (e.g., the genome, epigenome, transcriptome and proteome), which are biologically related and can each influence the behavior of the tumor. Thus, when studying molecular signals in cancer it is often necessary to consider data from multiple omics sources at once. This and other applications have motivated a very active research area in statistical methods for multi-omics integration.

A common task in multi-omics applications is to jointly characterize the molecular heterogeneity of the samples. Several multi-omics methods have been developed for this purpose, which can be broadly categorized by (1) clustering methods that identify molecularly distinct subtypes of the samples (Huo and Tseng, 2017; Lock and Dunson, 2013; Gabasova et al., 2017), (2) factorization methods that identify continuous lower-dimensional patterns of molecular variability (Lock et al., 2013; Argelaguet et al., 2018; Gaynanova and Li, 2019), or methods that combine aspects of (1) and (2) (Shen et al., 2013; Mo et al., 2017; Hellton and Thoresen, 2016). These extend classical approaches, such as (1) k-means clustering and (2) principal components analysis, to the multi-omics context, allowing the exploration of heterogeneity that is shared across the different 'omics sources while accounting for their differences.

Several multi-omics analyses have been performed on the TCGA data, including flagship publications for each type of cancer (e.g., see TCGA Research Network et al. (2012, 2014);

Verhaak et al. (2010)). These have revealed striking molecular heterogeneity within each classical type of cancer, which is often clinically relevant. However, restricting an analysis to a particular type of cancer sacrifices power to detect important genomic changes that are present across more than one cancer type. In 2013 TCGA began the Pan-Cancer Analysis Project, motivated by the observation that “cancers of disparate organs reveal many shared features, and, conversely, cancers from the same organ are often quite distinct” (Weinstein et al., 2013). Subsequently, several pan-cancer studies have identified important shared molecular alterations for somatic mutations (Kandoth et al., 2013), copy number (Zack et al., 2013), mRNA (Hoadley et al., 2014), and protein abundance (Akbari et al., 2014). However, a multi-omics analysis found that pan-cancer molecular heterogeneity is largely dominated by cell-of-origin and other factors that define the classical cancer types (Hoadley et al., 2018).

In this study we do not focus on baseline molecular differences between the cancer types. Rather, we focus on whether patterns of variability within each cancer type are shared across cancer types, i.e., whether multi-omic molecular profiles that drive heterogeneity in one type of cancer also drive heterogeneity in other cancers. Systematic investigations of heterogeneity in a pan-omics and pan-cancer context are presently limited by a lack of principled and computationally feasible statistical approaches for the comprehensive analysis of such data. In particular, the data take the form of *bidimensionally linked matrices*, i.e., multiple large matrices that may share row sets (here, defined by the omics platforms) or column sets (here, defined by the cancer types); this is illustrated in Figure 1 and the formal framework is described in Section 2.

In this article we propose a flexible approach to the simultaneous factorization and decomposition of variation across bidimensionally linked matrices, BIDIFAC+. This decomposes variation into a series of low-rank components that may be shared across any number

of row sets (e.g., omics platforms) or column sets (e.g., cancer types). Our approach builds on a growing literature for the factorization and decomposition of linked matrices, which we review in Section 3. Crucially, previous methods have primarily focused on multiple matrices that are linked in one dimension (rows or columns) only.




Figure 1: Bidimensional integration of pan-omics pan-cancer data.

2 Formal framework and notation

Here we introduce our framework and notation for pan-omics pan-cancer data. Let $\mathbf{X}_{ij} : M_i \times N_j$ denote the data matrix for omics data source i and sample set (i.e., cancer type) j for $j = 1, \dots, J$ and $i = 1, \dots, I$. Columns of \mathbf{X}_{ij} represent samples, and rows represent variables (e.g., genes, miRNAs, proteins). The sample sets of size $\mathbf{N} = [N_1, \dots, N_J]$ are consistent across each omics source, and the features measured for each omics source $\mathbf{M} = [M_1, \dots, M_I]$ are consistent across sample sets. As illustrated in Figure 1, the collection of available data can be represented as a single data matrix $\mathbf{X}_{..} : M \times N$ where $M = M_1 + \dots + M_I$ and $N = N_1 + \dots + N_J$, by horizontally and vertically concatenating its constituent blocks:

$$\mathbf{X}_{..} = \begin{bmatrix} \mathbf{X}_{11} & \mathbf{X}_{12} & \dots & \mathbf{X}_{1J} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{X}_{I1} & \mathbf{X}_{I2} & \dots & \mathbf{X}_{IJ} \end{bmatrix} \quad \text{where } \mathbf{X}_{ij} \text{ are } M_i \times N_j. \quad (1)$$

Similarly, $\mathbf{X}_{i\cdot}$ defines the concatenation of omics source i across cancer types and $\mathbf{X}_{\cdot j}$ defines the concatenation of cancer type j across omics sources:

$$\mathbf{X}_{i\cdot} = [\mathbf{X}_{i1} \dots \mathbf{X}_{iJ}] , \mathbf{X}_{\cdot j} = [\mathbf{X}'_{1j} \dots \mathbf{X}'_{Ij}]'$$

The notation $\mathbf{X}_{ij}[\cdot, n]$ defines the n 'th column of matrix ij , $\mathbf{X}_{ij}[m, \cdot]$ defines the m 'th row, and $\mathbf{X}_{ij}[m, n]$ defines the entry in row m and column n . In our context, the entries are all quantitative, continuous measurements; missing data are addressed in Section 9.

We will investigate shared or unique patterns of systematic variability (i.e., heterogeneity) among the constituent data blocks. We are not interested in baseline differences between the different omics platforms or sample sets, and so after routine preprocessing the data will be centered so that the mean of the entries within each data block, \mathbf{X}_{ij} , is

0. Moreover, to resolve the disparate scale of the data blocks, each block will be scaled to have comparable variability as described in Section 6.1.

In what follows, $\|\mathbf{A}\|_F$ denotes the Frobenius norm for any given matrix, so that $\|\mathbf{A}\|_F^2$ is the sum of squared entries in \mathbf{A} . The operator $\|\mathbf{A}\|_*$ denotes the nuclear norm of \mathbf{A} , which is given by the sum of the singular values in \mathbf{A} ; that is, if $\mathbf{A} : M \times N$ has ordered singular values $\mathbf{D}[1, 1], \mathbf{D}[2, 2], \dots$, then $\|\mathbf{A}\|_* = \sum_{r=1}^{\min(M, N)} \mathbf{D}[r, r]$.

3 Existing integrative factorization methods

There is now an extensive literature on the integrative factorization and decomposition of multiple linked datasets that share a common dimension. Much of this methodology is motivated by multi-omics integration, i.e., *vertical integration* of multiple matrices $\{\mathbf{X}_{11}, \mathbf{X}_{21}, \dots, \mathbf{X}_{M1}\}$ with shared columns in the setting of Section 2. For example, the Joint and Individual Variation Explained (JIVE) method (Lock et al., 2013; O’Connell and Lock, 2016) decomposes variation into *joint* components that are shared among multiple omics platforms and *individual* components that only explain substantial variability in one platform. This distinction simplifies interpretation, and also improves accuracy in recovering underlying signals. Accuracy improves because structured individual variation can interfere with finding important joint signal, just as joint structure can obscure important signal that is individual to a data source. The factorized JIVE decomposition is

$$\mathbf{X}_{i1} = \mathbf{U}_i \mathbf{V}^T + \mathbf{U}_i^* \mathbf{V}_i^T + \mathbf{E}_i \quad \text{for } i = 1, \dots, I. \quad (2)$$

Joint structure is represented by the common score matrix $\mathbf{V} : N_1 \times R$, which summarize patterns in the samples that explain variability across multiple omics platforms. The loading matrices $\mathbf{U}_i : M_i \times R$ indicate how these joint scores are expressed in the rows

(variables) of platform i . The score matrices $\mathbf{V}_i : N_1 \times R_i$ summarize sample patterns specific to platform i , with loadings \mathbf{U}_i^* . Model (2) can be equivalently represented as a sum of low-rank matrices

$$\mathbf{X}_{\cdot 1} = \mathbf{S}_{\cdot}^{(0)} + \sum_{i=1}^I \mathbf{S}_{\cdot}^{(i)} + \mathbf{E}. \quad (3)$$

where $\mathbf{S}_{\cdot}^{(0)} = \mathbf{U} \cdot \mathbf{V}^T$ is of rank R and $\mathbf{S}_{\cdot}^{(k)} = [\mathbf{S}_1^{(i)'} \dots \mathbf{S}_I^{(i)'}]'$ is the matrix of rank R_k given by the individual structure for platform k and zeros elsewhere:

$$\mathbf{S}_{i'}^{(i)} = \begin{cases} \mathbf{0}_{M_{i'} \times N} & \text{if } i' \neq i \\ \mathbf{U}_{i'}^* \mathbf{V}_i & \text{if } i' = i. \end{cases}$$

Several other methods result in a factorized decomposition similar to that in (2) and (3), including approaches that allow for different distributional assumptions on the constituent matrices (Li et al., 2018; Zhu et al., 2019), non-negative factorization (Yang and Michailidis, 2016), and the incorporation of covariates (Li and Jung, 2017). The Structural Learning and Integrative Decomposition (SLIDE) method (Gaynanova and Li, 2019) allows for a more flexible decomposition in which some components are only partially shared across a subset of the constituent data matrices. SLIDE extends model (3) to the more general decomposition

$$\mathbf{X}_{\cdot 1} = \sum_{k=1}^K \mathbf{S}_{\cdot}^{(k)} + \mathbf{E}. \quad (4)$$

where $\mathbf{S}_{\cdot}^{(k)} = [\mathbf{S}_1^{(k)'} \dots \mathbf{S}_I^{(k)'}]'$ is a low-rank matrix with non-zero values for some subset of the sources that is identified by a binary matrix $\mathbf{R} : I \times K$: and

$$\mathbf{S}_i^{(k)} = \begin{cases} \mathbf{0}_{M_i \times N} & \text{if } \mathbf{R}[i, k] = 0 \\ \mathbf{U}_i^{(k)} \mathbf{V}^{(k)T} & \text{if } \mathbf{R}[i, k] = 1 \end{cases}$$

Here, $\mathbf{V}^{(k)}$ gives scores that explain variability for only those patterns for the omics sources identified by $\mathbf{R}[\bullet, k]$.

The BIDIFAC approach (Park and Lock, 2019) is designed for the decomposition of bidimensionally linked matrices as in (1). Its low-rank factorization can be viewed as an extension of that for JIVE, decomposing variation into structure that is shared globally (G), across rows (Row), across columns (Col), or individual to the constituent matrices (Ind). Following (3) for JIVE and (4) for SLIDE, its full decomposition can be expressed as

$$\mathbf{X}_{..} = \mathbf{S}_{..}^{(G)} + \sum_{i=1}^I \mathbf{S}_{..}^{(i, \text{Row})} + \sum_{j=1}^J \mathbf{S}_{..}^{(j, \text{Col})} + \sum_{i=1}^I \sum_{j=1}^J \mathbf{S}_{..}^{(i, j, \text{Ind})} + \mathbf{E}_{..} \quad (5)$$

where $\mathbf{S}_{ij}^{(G)} = \mathbf{U}_i^{(G)} \mathbf{V}_j^{(G)T}$,

$$\mathbf{S}_{i'j'}^{(i, \text{Row})} = \begin{cases} \mathbf{0}_{M_i \times N_j} & \text{if } i' \neq i \\ \mathbf{U}_i^{(i, \text{Row})} \mathbf{V}_j^{(i, \text{Row})T} & \text{if } i' = i \end{cases}, \quad \mathbf{S}_{i'j'}^{(j, \text{Col})} = \begin{cases} \mathbf{0}_{M_i \times N_j} & \text{if } j' \neq j \\ \mathbf{U}_i^{(j, \text{Col})} \mathbf{V}_j^{(j, \text{Col})T} & \text{if } j' = j \end{cases},$$

and

$$\mathbf{S}_{i'j'}^{(i, j, \text{Ind})} = \begin{cases} \mathbf{0}_{M_i \times N_j} & \text{if } i' \neq i \text{ or } j' \neq j \\ \mathbf{U}_i^{(i, j, \text{Ind})} \mathbf{V}_j^{(i, j, \text{Ind})T} & \text{if } i' = i \text{ and } j' = j \end{cases}.$$

4 Proposed model

We consider a flexible factorization of bidimensionally linked data that combines aspects of the BIDIFAC and SLIDE models. Our full decomposition can be expressed as

$$\mathbf{X}_{..} = \sum_{k=1}^K \mathbf{S}_{..}^{(k)} + \mathbf{E}_{..}, \quad (6)$$

where

$$\mathbf{S}_{..}^{(k)} = \begin{bmatrix} \mathbf{S}_{11}^{(k)} & \mathbf{S}_{12}^{(k)} & \dots & \mathbf{S}_{1J}^{(k)} \\ \vdots & \vdots & \vdots & \vdots \\ \mathbf{S}_{I1}^{(k)} & \mathbf{S}_{I2}^{(k)} & \dots & \mathbf{S}_{IJ}^{(k)} \end{bmatrix}$$

and the presence of each $\mathbf{S}_{ij}^{(k)}$ is determined by a binary matrix of row indicators $\mathbf{R} : I \times K$ and column indicators $\mathbf{C} : J \times K$:

$$\mathbf{S}_{ij}^{(k)} = \begin{cases} \mathbf{0}_{M_i \times N_j} & \text{if } \mathbf{R}[i, k] = 0 \text{ or } \mathbf{C}[j, k] = 0 \\ \mathbf{U}_i^{(k)} \mathbf{V}_j^{(k)T} & \text{if } \mathbf{R}[i, k] = 1 \text{ and } \mathbf{C}[j, k] = 1 \end{cases}.$$

Each $\mathbf{S}_{..}^{(k)}$ gives a low-rank *module* that explains variability within the submatrix defined by the omics platforms identified by $\mathbf{R}[\cdot, k]$ and the cancer types identified by $\mathbf{C}[\cdot, k]$. There are in total $(2^I - 1)(2^J - 1)$ such submatrices, so by default we set $K = (2^I - 1)(2^J - 1)$ and let \mathbf{R} and \mathbf{C} enumerate all possible modules (see Appendix B). The SLIDE decomposition (4) is a special case when $J = 1$ or $I = 1$ (i.e., unidimensional integration); the BIDIFAC model (5) is a special case where each column of \mathbf{R} and \mathbf{C} contains either entirely 1's (i.e., all rows or columns included) or just one 1 (i.e., just one row set or column set included). The matrix $\mathbf{E}_{..}$ is an error matrix, whose entries are assumed to be sub-Gaussian with mean 0 and variance 1.

Let the rank of each module be $\text{rank}(\mathbf{S}_{..}^{(k)}) = R_k$, so that the dimensions of the non-zero components of the factorization are $\mathbf{U}_i^{(k)} : M_i \times R_k$ and $\mathbf{V}_j^{(k)} : N_j \times R_k$. The r 'th component of the k 'th module gives a (potentially multi-omic) molecular profile $\{\mathbf{U}_i^{(k)}[\cdot, r] : \mathbf{R}[i, k] = 1\}$ that explains variability within those cancer types defined by $\mathbf{C}[\cdot, k]$ with corresponding sample scores $\{\mathbf{V}_j^{(k)}[r, \cdot] : \mathbf{C}[i, k] = 1\}$.

5 Objective function

To estimate model (6), we minimize a least squares criterion with a structured nuclear norm penalty:

$$\operatorname{argmin}_{\{\mathbf{S}_{\cdot\cdot}^{(k)}\}_{k=1}^K} \frac{1}{2} \|\mathbf{X}_{\cdot\cdot} - \sum_{k=1}^K \mathbf{S}_{\cdot\cdot}^{(k)}\|_F^2 + \sum_{k=1}^K \lambda_k \|\mathbf{S}_{\cdot\cdot}^{(k)}\|_* \quad (7)$$

subject to $\mathbf{S}_{ij}^{(k)} = \mathbf{0}_{M_i \times N_j}$ if $\mathbf{R}[i, k] = 0$ or $\mathbf{C}[j, k] = 0$. The choice of the penalty parameters $\{\lambda_k\}_{k=1}^K$ is critical, and must satisfy the conditions of Proposition 8 to allow for non-zero estimation of each module.

Proposition 1. *Under objective (7), the following are necessary to allow for each $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}$ to be non-zero*

1. *If for $k' \neq k$ the rows and columns of module k' are contained within those for module k , $\mathbf{R}[i, k] - \mathbf{R}[i, k'] \geq 0 \forall i$ and $\mathbf{C}[j, k] - \mathbf{C}[j, k'] \geq 0 \forall j$, then $\lambda_k > \lambda_{k'}$.*
2. *If $\mathcal{I}_k \subset \{1, \dots, k-1, k+1, \dots, K\}$ is any subset of modules that together cover the rows and columns of module k , $\sum_{j \in \mathcal{I}_k} \mathbf{R}[\cdot, j] = r \cdot \mathbf{R}[\cdot, k]$ and $\sum_{j \in \mathcal{I}_k} \mathbf{C}[\cdot, j] = c \cdot \mathbf{C}[\cdot, k]$ for positive integers r and c , then $\lambda_k < \sum_{j \in \mathcal{I}_k} \lambda_j$.*

We determine the λ_k 's by random matrix theory, motivated by two well-known results for a single matrix that we repeat here in Propositions 2 and 3.

Proposition 2. *(Mazumder et al., 2010) Let $\mathbf{U}\mathbf{D}\mathbf{V}^T$ be the SVD of a matrix \mathbf{X} . The approximation \mathbf{A} that minimizes*

$$\frac{1}{2} \|\mathbf{X} - \mathbf{A}\|_F^2 + \lambda \|\mathbf{A}\|_* \quad (8)$$

is $\mathbf{A} = \mathbf{U}\tilde{\mathbf{D}}\mathbf{V}^T$, where $\tilde{\mathbf{D}}$ is diagonal with entries $\tilde{\mathbf{D}}[r, r] = \max(\mathbf{D}[r, r] - \lambda, 0)$.

Proposition 3. (Rudelson and Vershynin, 2010) If $\mathbf{E} : M \times N$ is a matrix of independent entries with mean 0 and variance σ^2 , then $\sigma(\sqrt{M} + \sqrt{N})$ provides a tight upper bound on the largest singular value ($\mathbf{D}[1, 1]$) of \mathbf{E} .

Fixing $\lambda = \sigma(\sqrt{M} + \sqrt{N})$ is a reasonable choice for the matrix approximation task in (8), because it keeps only those components r whose signal is greater than that expected for independent error by Proposition 3: $\mathbf{D}[r, r] > \sigma(\sqrt{M} + \sqrt{N})$ (Shabalin and Nobel, 2013). In our context $\sigma = 1$, and thus we set $\lambda_k = \sqrt{\mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, k] \cdot \mathbf{N}}$, where $\mathbf{R}[\cdot, k] \cdot \mathbf{M} \times \mathbf{C}[\cdot, k] \cdot \mathbf{N}$ gives the dimensions of the non-zero sub-matrix for $\mathbf{S}_{\cdot\cdot}^{(k)}$:

$$\mathbf{R}[\cdot, k] \cdot \mathbf{M} = \sum_{i=1}^I M_i \mathbf{R}[i, k] \quad \text{and} \quad \mathbf{C}[\cdot, k] \cdot \mathbf{N} = \sum_{j=1}^J N_j \mathbf{C}[j, k].$$

Our choice of λ_k is motivated to distinguish signal from noise in module $\mathbf{S}_{\cdot\cdot}^{(k)}$, conditional on the other modules. Moreover, it is guaranteed to satisfy the necessary conditions in Proposition 8, which we establish in Proposition 4.

Proposition 4. Setting $\lambda_k = \sqrt{\mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, k] \cdot \mathbf{N}}$ in (7) satisfies the necessary conditions of Proposition 8.

A similarly motivated choice of penalty weights is used in the BIDIFAC method, which solves an equivalent objective under the restricted scenario where the columns of \mathbf{R} and \mathbf{C} are fixed and contain either entirely 1's (i.e., all rows or columns included) or just one 1 (i.e., just one row set or column set included). Thus, we call our more flexible approach BIDIFAC+.

It is often infeasible to explicitly consider each of the $K = (2^I - 1)(2^J - 1)$ possible modules in (7), and the solution is often sparse, with $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}$ for several k . Thus, in practice we also optimize over the row and column sets \mathbf{R} and \mathbf{C} for some smaller number

of modules $\tilde{K} \ll K$:

$$\operatorname{argmin}_{\mathbf{R}, \mathbf{C}, \{\mathbf{S}_{\cdot}^{(k)}\}_{k=1}^{\tilde{K}}} \frac{1}{2} \|\mathbf{X}_{\cdot\cdot} - \sum_{k=1}^{\tilde{K}} \mathbf{S}_{\cdot}^{(k)}\|_F^2 + \sum_{k=1}^{\tilde{K}} \left(\sqrt{\mathbf{M} \cdot \mathbf{R}[\cdot, k]} + \sqrt{\mathbf{N} \cdot \mathbf{C}[\cdot, k]} \right) \|\mathbf{S}_{\cdot}^{(k)}\|_*. \quad (9)$$

Note that if \tilde{K} is greater than the number of non-zero modules, then the solution to (9) is equivalent to the solution to (7) in which \mathbf{R} and \mathbf{C} are fixed and enumerate all possible modules. If \tilde{K} is not greater than the number of non-zero modules, then the solution to (9) can still be informative as the set of \tilde{K} modules that together give the best structural approximation via (7).

6 Estimation

6.1 Scaling

We center each dataset \mathbf{X}_{ij} to have mean 0, and scale each dataset to have residual variance $\operatorname{var}(\mathbf{E}_{ij})$ of approximately 1. Such scaling requires an estimate of the residual variance for each dataset. By default we use the median absolute deviation estimator $\hat{\sigma}_{MAD}^2$ of Gavish and Donoho (2017), which is motivated by random-matrix theory under the assumption that \mathbf{X}_{ij} is composed of low-rank structure and mean 0 independent noise of variance σ^2 . We estimate $\hat{\sigma}_{MAD}^2$ for the unscaled data $\mathbf{X}_{ij}^{\text{unscaled}}$, and set $\mathbf{X}_{ij} = \mathbf{X}_{ij}^{\text{unscaled}} / \hat{\sigma}_{MAD}$. An alternative approach is to scale each dataset to have overall variance 1, $\operatorname{var}(\mathbf{X}_{ij})=1$, which is more conservative because $\operatorname{var}(\mathbf{E}_{\cdot\cdot}) \leq \operatorname{var}(\mathbf{X}_{ij})$; thus, this approach results in relatively larger λ_k in the objective function and leads to sparser overall ranks.

6.2 Optimization algorithm

We estimate across all modules $k = 1, \dots, K$ simultaneously by iteratively optimizing the objectives in Section 5. First assume the row and column inclusions for each module, defined by \mathbf{R} and \mathbf{C} , are fixed as in objective (7). Then, to estimate $\mathbf{S}_{..}^{(k)}$ given the other modules $\{\mathbf{S}_{..}^{(k')}\}_{k' \neq k}$, we can apply the soft-singular value estimator in Proposition 2 to the residual matrix

$$\mathbf{X}_{..} - \sum_{k' \neq k} \mathbf{S}_{..}^{(k')}$$

on the submatrix defined by $\mathbf{R}[\cdot, k]$ and $\mathbf{C}[\cdot, k]$. In this way, we iteratively optimize (7) over the K modules $\{\mathbf{S}_{..}^{(k)}\}_{k=1}^K$ until convergence. If the row and column inclusions \mathbf{R} and \mathbf{C} are not predetermined, then we incorporate additional sub-steps to estimate the non-zero submatrix defined by $\mathbf{R}[\cdot, k]$ and $\mathbf{C}[\cdot, k]$ for each module to optimize 9. We use a dual forward-selection procedure to iteratively determine the optimal row-set $\mathbf{R}[\cdot, k]$ with columns $\mathbf{C}[\cdot, k]$ fixed, and the column-set $\mathbf{C}[\cdot, k]$ with rows $\mathbf{R}[\cdot, k]$ fixed, until convergence prior to estimating each $\mathbf{S}_{..}^{(k)}$. Further details and pseudocode for the algorithm are provided in Appendix A.

7 Identifiability

Here we consider the identifiability of the decomposition in (4) under the objective (7). To account for permutation invariance of the K modules, throughout this section we assume that \mathbf{R} and \mathbf{C} are fixed and that they enumerate all of the $K = (2^I - 1)(2^J - 1)$ possible modules. Without loss of generality, we fix \mathbf{R} and \mathbf{C} as in Appendix B. Then, let $\mathbb{S}_{\mathbf{X}}$ be

the set of possible decompositions that yield a given approximation $\hat{\mathbf{X}}_{\bullet}$:

$$\mathbb{S}_{\hat{\mathbf{X}}} = \left\{ \{\mathbf{S}_{\bullet}^{(k)}\}_{k=1}^K \mid \hat{\mathbf{X}}_{\bullet} = \sum_{k=1}^K \mathbf{S}_{\bullet}^{(k)} \right\}.$$

If either $I > 1$ or $J > 1$ then the cardinality of $\mathbb{S}_{\hat{\mathbf{X}}}$ is infinite, i.e., there are an infinite number of ways to decompose $\hat{\mathbf{X}}_{\bullet}$. Thus, model (4) is clearly not identifiable in general, even in the no-noise case $\mathbf{E}_{\bullet} = \mathbf{0}$. However, optimizing the structured nuclear norm penalty in (7) may uniquely identify the decomposition; let $f_{\text{pen}}(\cdot)$ give this penalty:

$$f_{\text{pen}}(\{\mathbf{S}_{\bullet}^{(k)}\}_{k=1}^K) = \sum_{k=1}^K \left(\sqrt{\mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, k] \cdot \mathbf{N}} \right) \|\mathbf{S}_{\bullet}^{(k)}\|_{*}.$$

Proposition 5, gives an equivalence of the left and right singular vectors for any two decompositions that minimize $f_{\text{pen}}(\cdot)$.

Proposition 5. *Take two decompositions $\{\hat{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and $\{\tilde{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$, and assume that both minimize the structured nuclear norm penalty:*

$$f_{\text{pen}}(\{\hat{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K) = f_{\text{pen}}(\{\tilde{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K) = \min_{\mathbb{S}_{\hat{\mathbf{X}}}} f_{\text{pen}}(\{\mathbf{S}_{\bullet}^{(k)}\}_{k=1}^K).$$

Then, $\hat{\mathbf{S}}_{\bullet}^{(k)} = \mathbf{U}_{\bullet}^{(k)} \hat{\mathbf{D}}^{(k)} \mathbf{V}_{\bullet}^{(k)T}$ and $\tilde{\mathbf{S}}_{\bullet}^{(k)} = \mathbf{U}_{\bullet}^{(k)} \tilde{\mathbf{D}}^{(k)} \mathbf{V}_{\bullet}^{(k)T}$ where $\mathbf{U}_{\bullet}^{(k)} : M \times R_k$ and $\mathbf{V}_{\bullet}^{(k)} : N \times R_k$ have orthonormal columns, and $\hat{\mathbf{D}}^{(k)}$ and $\tilde{\mathbf{D}}^{(k)}$ are diagonal.

The proof of Proposition 5 uses two novel lemmas (see Appendix C): one establishing that $\hat{\mathbf{S}}_{\bullet}^{(k)}$ and $\tilde{\mathbf{S}}_{\bullet}^{(k)}$ must be additive in the nuclear norm, $\|\hat{\mathbf{S}}_{\bullet}^{(k)} + \tilde{\mathbf{S}}_{\bullet}^{(k)}\|_{*} = \|\hat{\mathbf{S}}_{\bullet}^{(k)}\|_{*} + \|\tilde{\mathbf{S}}_{\bullet}^{(k)}\|_{*}$, and a general result establishing that any two matrices that are additive in the nuclear norm must have the equivalence in Proposition 5.

Proposition 5 implies that left or right singular vectors of $\hat{\mathbf{S}}_{\bullet}^{(k)}$ ($\hat{\mathbf{D}}^{(k)}[r, r] > 0$) are either shared with $\tilde{\mathbf{S}}_{\bullet}^{(k)}$ (if $\tilde{\mathbf{D}}^{(k)}[r, r] > 0$) or orthogonal to $\tilde{\mathbf{S}}_{\bullet}^{(k)}$ (if $\tilde{\mathbf{D}}^{(k)}[r, r] = 0$). For identifiability, one must establish that $\hat{\mathbf{D}}^{(k)}[r, r] = \tilde{\mathbf{D}}^{(k)}[r, r]$ for all k and r . Theorem 2 gives sufficient conditions for identifiability of the decomposition.

Theorem 1. Consider $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and let $\mathbf{U}_{\cdot}^{(k)}\hat{\mathbf{D}}^{(k)}\mathbf{V}_{\cdot}^{(k)T}$ give the SVD of $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}$ for $k = 1, \dots, K$. The following three properties uniquely identify $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$.

1. $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$ minimizes $f_{\text{pen}}(\cdot)$ over $\mathbb{S}_{\hat{\mathbf{X}}}$,
2. $\{\hat{\mathbf{U}}_i^{(k)}[\cdot, r] : \mathbf{R}[i, k] = 1 \text{ and } \hat{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $i = 1, \dots, I$,
3. $\{\hat{\mathbf{V}}_j^{(k)}[\cdot, r] : \mathbf{C}[j, k] = 1 \text{ and } \hat{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $j = 1, \dots, J$.

The linear independence conditions (2. and 3.) are in general not sufficient for identifiability, and several related integrative factorization methods achieve identifiability via stronger orthogonality conditions across the terms of the decomposition (Lock et al., 2013; Gaynanova and Li, 2019). Theorem 2 implies that orthogonality is not necessary under the penalty $f_{\text{pen}}(\cdot)$. Conditions 2. and 3. are straightforward to verify for any $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$, and they will generally hold whenever the ranks in the estimated factorization are small relative to the dimensions $\{M_i\}_{i=1}^I$ and $\{N_j\}_{j=1}^J$. Moreover, the conditions of Theorem 2 are only sufficient for identifiability; there may be cases for which the minimizer of $f_{\text{pen}}(\cdot)$ is unique and the terms of its decomposition are not linearly independent.

8 Bayesian interpretation

Express the BIDIFAC+ model (6) in factorized form

$$\mathbf{X}_{\cdot\cdot} = \sum_{k=1}^K \mathbf{U}_{\cdot}^{(k)}\mathbf{V}_{\cdot}^{(k)T} + \mathbf{E}_{\cdot\cdot} \quad (10)$$

where

$$\mathbf{U}_{\cdot}^{(k)'} = [\mathbf{U}_1^{(k)'} \dots \mathbf{U}_I^{(k)'}]', \text{ with } \mathbf{U}_i^{(k)} = M_1 \times R_k \text{ and } \mathbf{U}_i^{(k)} = \mathbf{0}_{M_1 \times R_k} \text{ if } \mathbf{R}[i, k] = 1 \quad (11)$$

for all i and k , and

$$\mathbf{V}_i^{(k)} = [\mathbf{V}_1^{(k)} \cdots \mathbf{V}_J^{(k)}], \text{ with } \mathbf{V}_i^{(k)} = N_j \times R_k \text{ and } \mathbf{V}_i^{(k)} = \mathbf{0}_{N_j \times R_k} \text{ if } \mathbf{C}[j, k] = 1 \quad (12)$$

for all j and k . The structured nuclear norm objective (7) can also be represented by L_2 penalties on the factorization components $\mathbf{U}_i^{(k)}$ and $\mathbf{V}_j^{(k)}$. We formally state this equivalence in Proposition 6, which extends analogous results for a single matrix (Mazumder et al., 2010) and for the BIDIFAC framework (Park and Lock, 2019).

Proposition 6. *Fix \mathbf{R} and \mathbf{C} . Let $\{\hat{\mathbf{U}}_i^{(k)}\}_{k=1}^K$ and $\{\hat{\mathbf{V}}_j^{(k)}\}_{k=1}^K$ minimize*

$$\|\mathbf{X}_{..} - \sum_{k=1}^K \mathbf{U}_i^{(k)} \mathbf{V}_j^{(k)}\|_F^2 + \sum_{k=1}^K \lambda_k (\|\hat{\mathbf{U}}_i^{(k)}\|_F^2 + \|\hat{\mathbf{V}}_j^{(k)}\|_F^2) \quad (13)$$

with the restrictions (11) and (12). Then, $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$ solves (7), where $\hat{\mathbf{S}}_{..}^{(k)} = \hat{\mathbf{U}}_i^{(k)} \hat{\mathbf{V}}_j^{(k)T}$ for $k = 1, \dots, K$.

From (13), it is apparent that our objective gives the mode of a Bayesian posterior with normal priors on the errors and the factorization components, as stated in Proposition 7.

Proposition 7. *Let the entries of $\mathbf{E}_{..}$ be independent $\text{Normal}(0, 1)$, the entries of $\mathbf{U}_i^{(k)}$ be independent $\text{Normal}(0, \tau^2)$ if $\mathbf{R}[i, k] = 1$, and the entries of $\mathbf{V}_j^{(k)}$ be independent $\text{Normal}(0, \tau_k^2)$ if $\mathbf{C}[j, k] = 1$, where $\tau_k^2 = 1/\lambda_k$. Then, (13) is proportional to the joint likelihood*

$$p(\mathbf{X}_{..}, \{\mathbf{U}_i^{(k)}\}_{k=1}^K, \{\mathbf{V}_j^{(k)}\}_{k=1}^K \mid \mathbf{R}, \mathbf{C}).$$

9 Missing data imputation

The probabilistic formulation of the objective described in Section 8 motivates a modified Expectation-Maximization (EM)-algorithm approach to impute missing data. Let \mathcal{M} index

observations in the full dataset $\mathbf{X}_{..}$ that are unobserved: $\mathcal{M} = \{(m, n) : \mathbf{X}_{..}[m, n] \text{ is missing}\}$.

Our iterative algorithm for missing data imputation proceeds as follows:

1. Initialize $\hat{\mathbf{X}}_{..}$ by

$$\hat{\mathbf{X}}_{..}[m, n] = \begin{cases} \mathbf{X}_{..}[m, n] & \text{if } (m, n) \notin \mathcal{M} \\ 0 & \text{if } (m, n) \in \mathcal{M} \end{cases}$$

2. M-step: Estimate $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$ by optimizing (7) for $\hat{\mathbf{X}}_{..}$.
3. E-step: Update $\hat{\mathbf{X}}_{..}$ by $\hat{\mathbf{X}}_{..}[m, n] = \sum_{k=1}^K \hat{\mathbf{S}}_{..}^{(k)}[m, n]$.
4. Repeat steps 2. and 3. until convergence.

Analogous approaches to imputation for other low-rank factorization techniques have been proposed (Kurucz et al., 2007; O’Connell and Lock, 2017; Park and Lock, 2019). Crucially for our context, the method can be used to impute data that may be missing from an entire column or an entire row of each \mathbf{X}_{ij} .

10 Application to TCGA data

10.1 Data acquisition and preprocessing

Our data were curated for the TCGA Pan-Cancer Project and were used for the pan-cancer clustering analysis described in Hoadley et al. (2018). We used data from four ($I = 4$) omics sources: (1) batch corrected RNA-Seq data capturing (mRNA) expression for 20531 genes, (2) batch corrected miRNA-Seq data capturing expression for 743 miRNAs, (3) between-platform normalized data from the Illumina 27K and 450K platforms capturing DNA methylation levels for 22601 CpG sites, and (4) batch-corrected reverse-phase protein array

data capturing abundance for 198 proteins. These data are available for download at <https://gdc.cancer.gov/about-data/publications/PanCan-CellOfOrigin> [accessed 11/19/2019]. We consider data for $N = 6,973$ tumor samples from different individuals with all four omics sources available; these tumor samples represent $J = 29$ different cancer types, listed in Table 1.

Table 1: TCGA acronyms for the 29 different cancer types considered.

Acronym	Cancer type	Acronym	Cancer type
ACC	Adrenocortical carcinoma	BLCA	Bladder urothelial carcinoma
BRCA	Breast invasive carcinoma	CESC	Cervical carcinoma
CHOL	Cholangiocarcinoma	CORE	Colorectal adenocarcinoma
DLBC	Diffuse large B-cell lymphoma	ESCA	Esophageal carcinoma
HNSC	Head/neck squamous cell	KICH	Kidney chromophobe
KIRC	Kidney renal clear cell	KIRP	Kidney renal papillary cell
LGG	Brain lower grade glioma	LIHC	Liver hepatocellular carcinoma
LUAD	Lung adenocarcinoma	LUSC	Lung squamous cell carcinoma
MESO	Mesothelioma	OV	Ovarian cancer
PAAD	Pancreatic adenocarcinoma	PCPG	Pheochromocytoma and paraganglioma
PRAD	Prostate adenocarcinoma	SARC	Sarcoma
SKCM	Skin cutaneous melanoma	STAD	Stomach adenocarcinoma
TGCT	Testicular germ cell tumors	THCA	Thyroid carcinoma
THYM	Thymoma	UCEC	Uterine corpus endometrial carcinoma
UCS	Uterine carcinosarcoma		

Table 2: Cancer types and sources for the first 15 modules, ordered by variation explained.

Module	Cancer types	Omics sources
1	All cancers	mRNA miRNA Meth Protein
2	All cancers	miRNA
3	BLCA BRCA CESC CHOL CORE DLBC ESCA HNSC LIHC LUAD LUSC OV PAAD PRAD SKCM STAD TGCT UCEC UCS	Meth
4	ACC BLCA CHOL CORE DLBC ESCA HNSC KICH KIRC KIRP LGG LIHC LUAD LUSC MESO PAAD PCPG SARC SKCM STAD THCA THYM	mRNA Meth
5	All cancers	mRNA
6	BRCA	mRNA miRNA Meth Protein
7	LGG	mRNA miRNA Protein
8	All cancers *but* LGG	Protein
9	THCA	mRNA miRNA Protein
10	All cancers *but* LGG and TGCT	miRNA
11	CHOL KIRC KIRP LIHC	mRNA miRNA Meth Protein
12	LGG	Meth
13	BLCA CESC CORE ESCA HNSC LUSC SARC STAD	mRNA miRNA Meth Protein
14	KICH KIRC KIRP	mRNA miRNA Protein
15	BLCA BRCA CESC CHOL ESCA HNSC LUAD LUSC PAAD PRAD SKCM STAD TGCT UCEC UCS	mRNA miRNA

We log-transformed the counts for the RNA-Seq and miRNA-Seq sources. To remove baseline differences between cancer types, we center each data source to have mean 0 across all rows for each cancer type:

$$\text{mean}(\mathbf{X}_{ij}[m, \cdot]) = 0 \text{ for all } i, j, m.$$

We filter to the 1000 genes and the 1000 methylation CpG probes that have the highest standard deviation after centering, leaving $M_1 = 1000$ genes, $M_2 = 743$ miRNAs, $M_3 = 1000$ CpGs, and $M_4 = 198$ proteins. Lastly, to account for differences in scale, we standardize so that each variable has standard deviation 1:

$$\text{SD}(\mathbf{X}_i[m, \cdot]) = 1 \text{ for all } i, m.$$

10.2 Factorization results

We apply the BIDIFAC+ method to the complete-case data with $I = 4$ omics sources and $J = 29$ cancer types. We simultaneously estimate a maximum of $K = 50$ low-rank modules; all modules are non-zero, but the variation explained by the smaller modules are negligible. Figure 10.2 gives the total variance explained by each module, $\|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_F^2$, for $k = 1, \dots, 50$ in decreasing order. The top 15 modules ordered by total variance explained are given in Table 2, and all 50 modules are given in the supplemental spreadsheet at http://www.ericfrazierlock.com/BIDIFAC_modules.xlsx. The first module explains global variation, with all cancer types and all omics sources included. Other modules that explain substantial variability across all or almost all cancer types are specific to each omics source: miRNA (Module 2), methylation (Module 3), gene expression (Module 5) and Protein (Module 8).

The module that explains the fourth most variation (Module 4) identifies structure in the genes and DNA methylation that explains variation in 22 of the 29 cancer types; we focus




Figure 2: Total sum of squared entries in each of the 50 modules.

on this module as an illustrative example. The cancer types *not* included in Module 4 are BRCA (breast), CESC (cervical), OV (ovarian), PRAD (prostate), TGCT (testicular), UCEC (uterine endometrial), and UCS (uterine). Interestingly, all tumor types that were excluded were cancers specific to either males or females (or heavily skewed in BRCA); while cancer types included have both sexes. Figure 3 shows that Module 4 is indeed dominated by a single component that corresponds to molecular differences between the sexes. The gene loadings for this component are negligible except for those on the Y chromosome and two genes on the X chromosome that are responsible for X-inactivation in females: *XIST* and *TSIX*; the methylation loadings are negligible except for those in the X chromosome. These results are an intuitive illustration of the method, revealing a multi-omic molecular

signal that explains heterogeneity in some cancer types, but not all cancer types (only those that have both males and females).

The module that explains the sixth most variation (Module 6) identifies structure across all four omics sources that explains variation in the breast cancer (BRCA) samples only. Figure 4 shows that the first two components in this module are driven primarily by distinctions between the PAM50 molecular subtypes for BRCA (TCGA Research Network et al., 2012). Thus, our analysis suggests that molecular signals that distinguish these subtypes are present across all four omics sources, but that these signals do not explain substantial variation within any other type of cancer considered.

Several other modules explain variability in just one type of type of cancer, including LGG (Module 7: mRNA, miRNA and Protein), THCA (Module 9), UCEC (Module 16), and PRAD (Modules 18 and 19). Module 12, which is specific to LGG methylation, reveals distinct clustering by mutation status of the *IDH* genes (see Figure 10.2). IDH mutations have been shown to lead to a distinct CpG-island hyper-methylated phenotype (Noushmehr et al., 2010). Other modules explain variability in multiple cancer types that share similarities regarding their origin or histology. For example, Module 14 explains variability within the three kidney cancers (KICH, KIRC, and KIRP), and digestive and gastrointestinal cancers (CORE, ESCA, PAAD, STAD) are represented in Modules 25 (methylation) and 28 (mRNA).

10.3 Missing data imputation

To assess the accuracy of missing data imputation using BIDIFAC+, we hold-out observed entries, rows, and columns of each dataset in the pan-omics pan-cancer and impute them using the approach in Section 9. We randomly set 100 columns (samples) to missing for each of the 4 omic platforms, and we randomly set 100 rows (features) to missing for

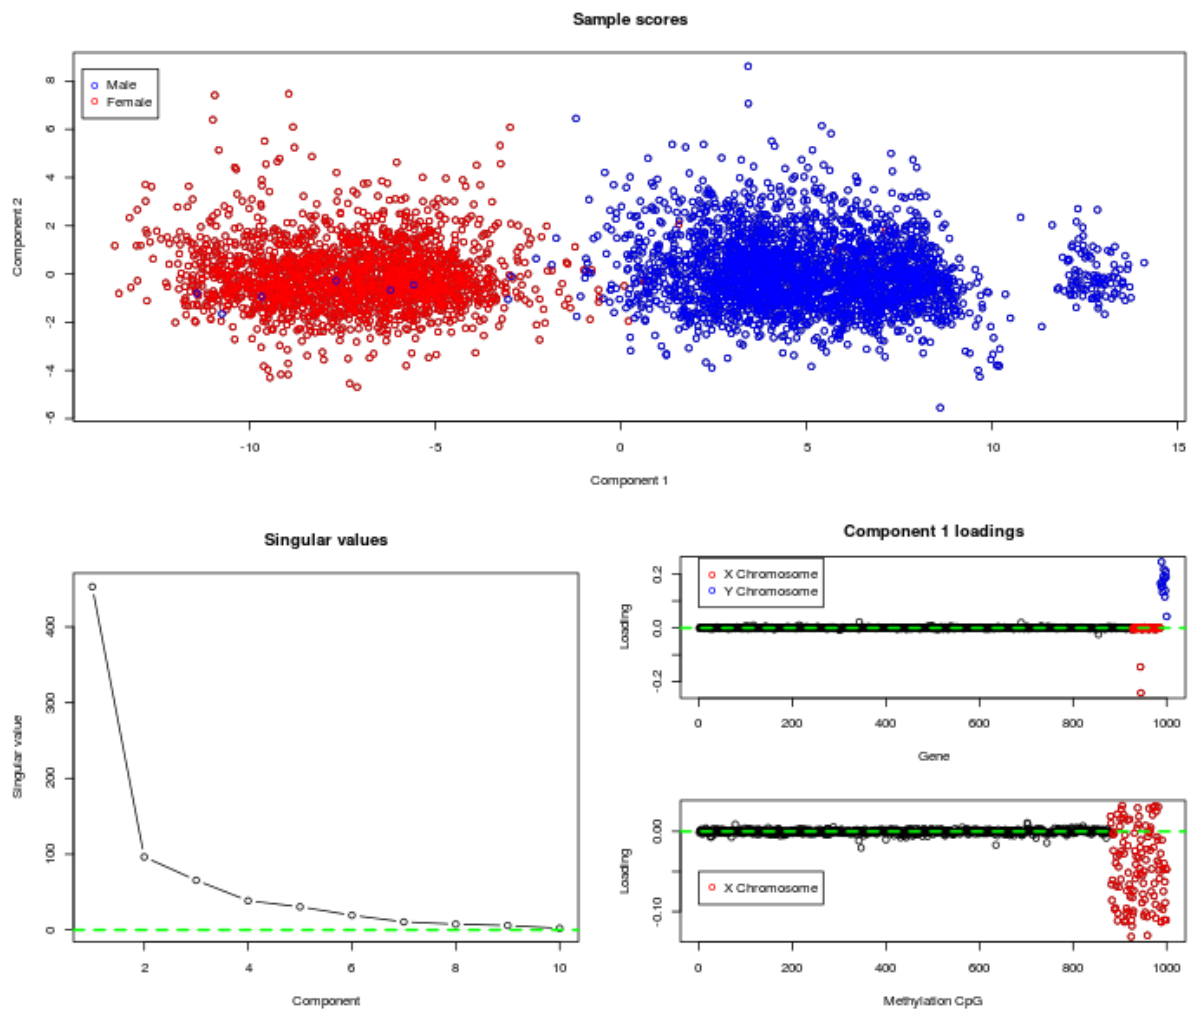


Figure 3: Sample scores (top), scree plot (bottom left), and loadings on genes and methylation CpGs (bottom right) for the first component of Module 4. This includes 22 cancer types with samples from both sexes, and the module is dominated by molecular signals that distinguish males from females.

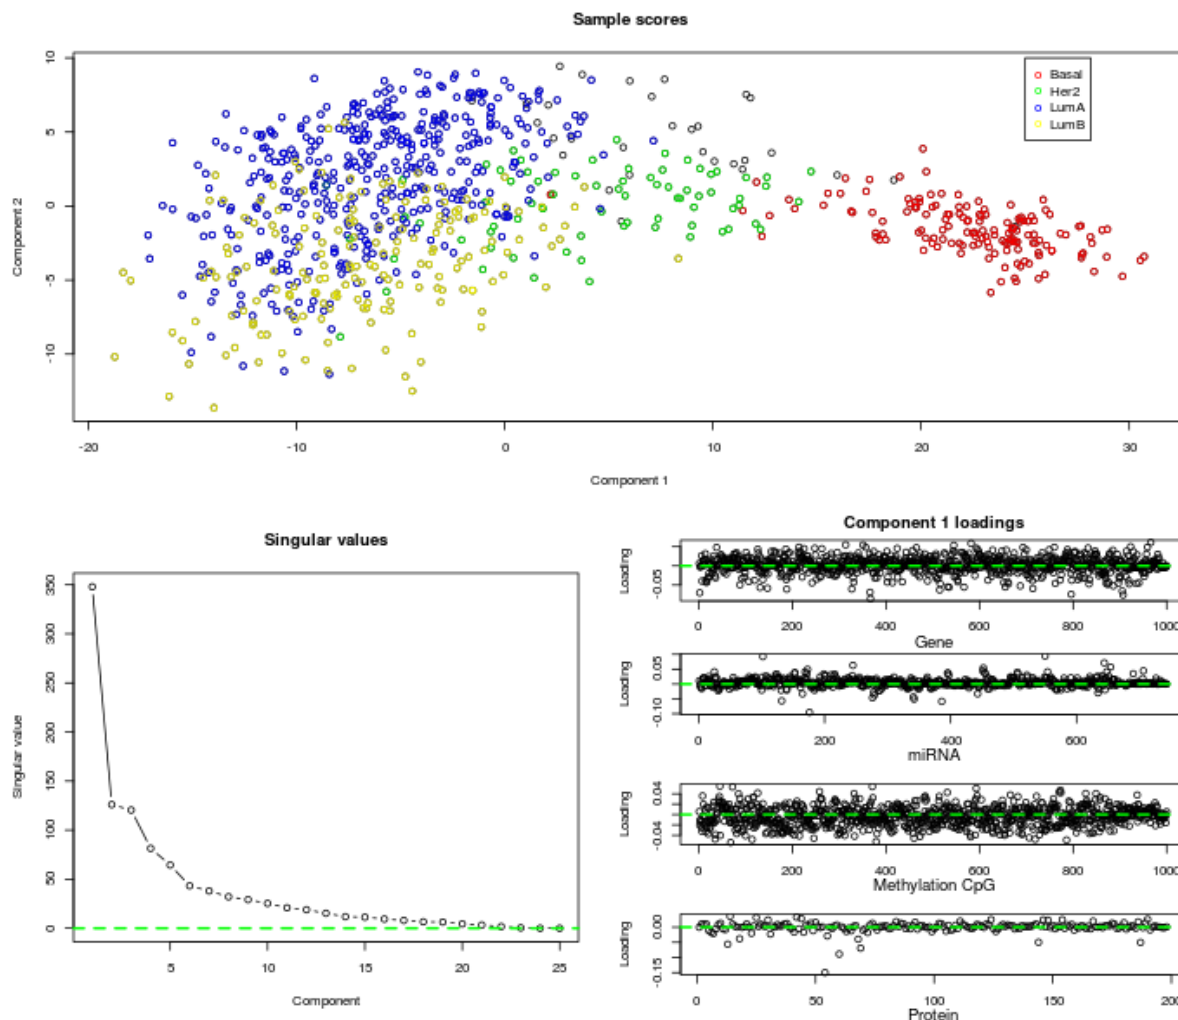


Figure 4: Sample scores (top), scree plot (bottom left), and loadings for all four omics platforms for Module 6. This module includes only breast (BRCA) tumor samples, and it is dominated by molecular signals that distinguish the PAM50 subtypes.

each of the 29 cancer types. We then randomly set 5000 of the values remaining in the joint matrix $\mathbf{X}_{..}$ to missing. We impute missing values using BIDIFAC+ as described in




Figure 5: Scores for the first two components of Module 12 (LGG; methylation), colored by IDH mutation status.

Section 9, and for comparison we use an analogous approach to imputation using four other low-rank factorizations: (1) soft-threshold (nuclear norm) SVD of the joint matrix \mathbf{X}_* , (2) soft-threshold SVD of each matrix \mathbf{X}_{ij} separately, (3) hard-threshold SVD (SVD approximation using the first R singular values) of \mathbf{X}_* , (4) hard-threshold SVD of each \mathbf{X}_{ij} separately. For the soft-thresholding SVD methods, the penalty factor is estimated by random matrix theory as in Section 5. For the hard-thresholding methods the ranks are determined by cross-validation by minimizing imputation error on an additional held-out

cohort of the same size.

We consider the imputation error under the different methods, broken down by (1) observed values, (2) values that are missing but have the rest of their row and column present (entrywise missing), (3) values that are missing their entire row, (4) values that are missing their entire column, and (5) values that are missing both their row and their column. For a given set of values \mathcal{M} , we compute the relative squared error as

$$\text{RSE} = \frac{\sum_{(m,n) \in \mathcal{M}} (\mathbf{X}_{..}[m, n] - \hat{\mathbf{X}}_{..}[m, n])^2}{\sum_{(m,n) \in \mathcal{M}} \mathbf{X}_{..}[m, n]^2},$$

where $\hat{\mathbf{X}}_{..}$ is the structural approximation resulting from the given method. Table 3 gives the RSE for each method and for each missing condition. Imputation by BIDIFAC+ outperforms the other methods for each type of missingness, illustrating the advantages of decomposing joint and individual structures. The hard-thresholding approaches have much less error for the observed data than for the missing data, due to over-fitting of the signal.

Table 3: Imputation RSE under different approaches and different types of missingness.

Method	Observed	Entrywise	Row	Column	Both
BIDIFAC+	0.510	0.558	0.670	0.807	0.881
Soft-SVD (joint)	0.531	0.621	0.678	0.834	0.894
Soft-SVD (separate)	0.564	0.610	1.000	1.000	1.000
Hard-SVD (joint)	0.431	0.559	0.829	0.908	1.200
Hard-SVD (separate)	0.344	0.581	1.000	1.000	1.000

11 Simulation studies

11.1 Vertically linked simulations

We conduct a simulation study to assess the accuracy of the BIDIFAC+ decomposition in the context of vertical integration, where there is a single shared column set ($J = 1$). For all scenarios, we simulate data according to model (4) wherein the entries of the residual noise $\mathbf{E}_.$ are generated independently from a Normal(0, 1) distribution and the entries of each $\mathbf{U}_i^{(k)}$ and $\mathbf{V}^{(k)}$ are generated independently from a Normal(0, σ^2) distribution.

We first consider a scenario with $I = 3$ matrices, each of dimension 100×100 ($N = 100$ and $M_1 = M_2 = M_3 = 100$), with low-rank modules that are shared jointly, shared across each pair of matrices, and individual to each matrix:

$$\mathbf{R} = \begin{bmatrix} 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 1 \end{bmatrix}. \quad (14)$$

We consider a “low-rank” and a “high-rank” condition across three different signal-to-noise levels. For the low-rank condition, each of the seven modules has rank $R = 1$; for the high-rank condition, each module has rank $R = 5$. The variance of the factorized signal component, σ^2 is set to be $\sqrt{1/2}$, 1, or $\sqrt{10}$, so that the signal-to-noise ratio (s2n) of each components is 1/2, 1, or 10, respectively.

For each condition, we apply four approaches to uncover the underlying decomposition:

1. BIDIFAC+, with \mathbf{R} given by (14), as in the true generative model,
2. BIDIFAC+, with \mathbf{R} estimated,
3. SLIDE, with \mathbf{R} and the true ranks of each module ($R = 1$ or $R = 5$) provided,

4. SLIDE, with \mathbf{R} and the ranks of each module estimated via the default cross-validation scheme.

We use SLIDE as the basis of comparison with BIDIFAC+, because it is the only other method that is designed to recover each term in the decomposition and it generally outperforms other vertically linked decomposition methods (Gaynanova and Li, 2019; Park and Lock, 2019). For each case we compute the mean relative squared error (RSE) in recovering each module of the decomposition:

$$\text{RSE} = \frac{1}{K} \sum_{k=1}^K \frac{\|\mathbf{S}_{\cdot\cdot}^{(k)} - \hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_F^2}{\|\mathbf{S}_{\cdot\cdot}^{(k)}\|_F^2}. \quad (15)$$

The mean RSE for each condition and under each approach is shown in Table 4, broken down by the global module, pairwise modules, and individual modules. BIDIFAC+ generally outperforms or performs similarly to SLIDE, even when the true ranks are used for the SLIDE implementation (the ranks are never fixed for BIDIFAC+). An exception is when the ranks and s2n ratio are small (rank=1, s2n=0.5), where BIDIFAC+ tends to over-shrink the signal. BIDIFAC+ performs particularly well relative to SLIDE when the rank is large and s2n is high. One likely reason for this improvement is that the SLIDE model necessarily restricts the factorized components $\mathbf{U}_i^{(k)}$ and $\mathbf{V}^{(k)}$ to be mutually orthogonal, whereas BIDIFAC+ has no such constraint. This restriction can be limiting when decomposing generated signals that are independent but not orthogonal Park and Lock (2019). Moreover, when estimating the ranks the SLIDE model can drastically underperform relative to using the true ranks. The results for BIDIFAC+ when fixing the true modules \mathbf{R} vs. estimating \mathbf{R} are nearly identical; because all possible modules are present for this scenario, the two approaches are very similar despite subtle differences in the algorithms.

We consider another scenario with a larger number of matrices ($I = 10$), each of dimension 100×100 ($N = 100$, $M_1 = \dots = M_{10} = 100$) and sparsely distributed modules. We

generate 10 low rank modules out of $2^{10} - 1 = 1023$ possibilities, that are present on (1) X_{11} only, (2) X_{11} and X_{21} , (3) X_{11} , X_{21} , and X_{31} , etc. We again consider low-rank ($R = 1$) and high-rank ($R = 5$) scenarios for all modules, and three signal-to-noise levels 0.5, 1, and 10. The resulting mean RSE (15) over all modules, for each approach, is shown in Table 5. Here, BIDIFAC+ with fixed true \mathbf{R} generally performs better than estimating \mathbf{R} ; however, these gains are modest for most scenarios, suggesting the BIDIFAC+ generally does a good job of identifying which of the 1023 possible modules are non-zero.

11.2 Application-motivated simulation

Here we assess the recovery of the underlying structure and the accuracy of the decomposition into shared components for a bidimensionally linked scenario that reflects our motivating application in Section 10. We generate data by taking the estimated decomposition from Section 10.2 and adding independent noise to it. That is, we simulate

$$\tilde{\mathbf{X}}_{..} = \sum_{k=1}^{50} \alpha \hat{\mathbf{S}}_{..}^{(k)} + \tilde{\mathbf{E}}_{..}$$

where $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$ is the estimated decomposition from Section 10.2, the entries of $\tilde{\mathbf{E}}_{..}$ are independent Normal(0, 1), and $\alpha > 0$ is a parameter that controls the total signal-to-noise ratio. We consider three total signal-to-noise ratios, defined by

$$\text{s2n} = \text{var}\left(\sum_{k=1}^{50} \alpha \hat{\mathbf{S}}_{..}^{(k)}\right) / \text{var}(\tilde{\mathbf{E}}) = \text{var}\left(\sum_{k=1}^{50} \alpha \hat{\mathbf{S}}_{..}^{(k)}\right),$$

$\text{s2n} = 0.2, 0.5$, and 5. The scenario with $\text{s2n} = 0.5$ corresponds most closely to the real data, for which the ratio of the estimated signal variance over the residual variance is 0.552. For each scenario, we estimate the underlying decomposition using BIDIFAC+ with the true \mathbf{R} and \mathbf{C} fixed, and using BIDIFAC+ with estimated modules $\tilde{\mathbf{R}}$ and $\tilde{\mathbf{C}}$ and $\tilde{K} = 50$.

Table 4: Comparison of BIDFAC+ and SLIDE signal decomposition RSE ($I = 3$ sources).

Scenario	Structure	BIDFAC+		SLIDE	
		True \mathbf{R}	Estimated \mathbf{R}	True ranks	Estimated ranks
Rank=1, s2n=0.5	Global	0.130	0.130	0.120	0.120
	Pairwise	0.157	0.156	0.103	0.103
	Individual	0.197	0.197	0.118	0.118
Rank=1, s2n=1	Global	0.060	0.060	0.084	0.084
	Pairwise	0.068	0.068	0.053	0.053
	Individual	0.070	0.070	0.048	0.048
Rank=1, s2n=10	Global	0.010	0.010	0.035	3.65
	Pairwise	0.005	0.005	0.027	1.00
	Individual	0.008	0.008	0.037	0.689
Rank=5, s2n=0.5	Global	0.270	0.270	0.276	0.869
	Pairwise	0.268	0.268	0.263	0.460
	Individual	0.329	0.329	0.306	0.317
Rank=5, s2n=1	Global	0.123	0.123	0.232	1.320
	Pairwise	0.121	0.121	0.189	0.674
	Individual	0.148	0.148	0.241	0.485
Rank=5, s2n=10	Global	0.080	0.080	0.233	2.36
	Pairwise	0.060	0.060	0.189	0.917
	Individual	0.089	0.089	0.249	0.703

Table 5: Comparison of BIDFAC+ and SLIDE signal decomposition RSE ($I = 10$ sources).

		BIDIFAC+		SLIDE	
Ranks	s2n	True \mathbf{R}	Estimated \mathbf{R}	True ranks	Estimated ranks
1	0.5	0.150	0.150	0.116	0.116
1	1	0.076	0.078	0.105	0.105
1	10	0.032	0.025	0.060	0.060
5	0.5	0.297	0.320	0.402	0.685
5	1	0.177	0.189	0.324	0.603
5	10	0.167	0.245	0.347	0.347

In each case, we compute the RSE as follows

$$\text{RSE} = \frac{1}{50} \sum_{k=1}^{50} \frac{\|\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)} - \alpha \hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_F^2}{\|\alpha \mathbf{S}_{\cdot\cdot}^{(k)}\|_F^2} \quad (16)$$

When computing RSE, we permute the 50 modules so that $\tilde{\mathbf{R}}[\cdot, k] = \mathbf{R}[\cdot, k]$ and $\tilde{\mathbf{C}}[\cdot, k] = \mathbf{C}[\cdot, k]$ wherever possible, and set $\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}$ if $\tilde{\mathbf{R}}[\cdot, k] \neq \mathbf{R}[\cdot, k]$ and $\tilde{\mathbf{C}}[\cdot, k] \neq \mathbf{C}[\cdot, k]$. We also compute the relative overall signal recovery (ROSR) as

$$\text{ROSR} = \frac{\|\sum_{k=1}^K \hat{\mathbf{S}}_{\cdot\cdot}^{(k)} - \sum_{k=1}^K \alpha \hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_F^2}{\|\sum_{k=1}^K \alpha \mathbf{S}_{\cdot\cdot}^{(k)}\|_F^2} \quad (17)$$

The results are shown in Table 6, and demonstrate that the underlying decomposition is recovered reasonably well in most scenarios. However, the RSE for estimated modules is often substantially more than the RSE using the true modules, as the row and column sets defining the modules can be estimated incorrectly. Moreover, the overall signal recovery error (ROSR) is generally substantially less than the mean error in recovering each module (RSE), demonstrating how the decomposition can be estimated incorrectly even if the overall signal is estimated with high accuracy.

Table 6: Relative squared error of the decomposition (RSE) and relative overall signal recovery (ROSR) using BIDIFAC+ with known modules (\mathbf{R} and \mathbf{C}) and estimated modules ($\hat{\mathbf{R}}$ and $\hat{\mathbf{C}}$).

s2n	RSE(\mathbf{R}, \mathbf{C})	RSE ($\tilde{\mathbf{R}}, \tilde{\mathbf{C}}$)	ROSR (\mathbf{R}, \mathbf{C})	ROSR ($\tilde{\mathbf{R}}, \tilde{\mathbf{C}}$)
0.2	0.356	0.531	0.170	0.189
0.5	0.242	0.386	0.131	0.143
5	0.128	0.346	0.012	0.026

12 Discussion

The successful integration of multiple large sources of data is a pivotal challenge for many modern analysis tasks. While several general approaches have been developed, they largely do not apply to the context of bidimensionally linked matrices. BIDIFAC+ is a flexible approach for dimension reduction and decomposition of shared structures among bidimensionally linked matrices, which is competitive with alternative methods that integrate over a single dimension (rows or columns). Here we have focused primarily on the accuracy of the estimated decomposition and exploratory analysis of the results. BIDIFAC+ may also be used for other tasks, such as missing data imputation or as a dimension reduction step preceding statistical modeling (e.g., as in principal components regression). For these other tasks it is desirable to model statistical uncertainty, and fully Bayesian extensions that capture the full posterior distribution about the mode in Section 8 are potentially very useful. Moreover, while we have explored the identifiability of the decomposition under BIDIFAC+, it is worthwhile to establish conditions that are both necessary and sufficient for identifiability.

Our application to pan-omics pan-cancer data from TCGA revealed molecular patterns

that explain variability across all or almost all types of cancer, both across omics platforms and within each omics platform. However, it also revealed patterns several instances in which patterns are specific to one or a small subset of cancers, and these often show sharp distinctions of previously known molecular subtypes (e.g., for BRCA and LGG). Interestingly, BRCA was the only tumor type that showed up with all four platforms in a module. Together, they strongly separated the Basal-like molecular subtype from other subtypes of breast cancer. This mirrors the analysis of individual data types in TCGA Research Network et al. (2012). The LGG data also split by both histological groups and mutation status based on BIDFAC+, even though both were not included in the analysis. Module 7 included mRNA, miRNA, and protein and was predominantly driven by co-deletion of 1p/19q which is predominantly observed in oligodendrogliomas and is associated with better overall survival. This mirrors the previous TCGA work that showed that the LGG could be predominately split by 1p/19q deletion, IHD1 status (Module 12, for methylation) or TP53 mutation status (TCGA Research Network, 2015).

Availability

R code to perform BIDIFAC+ and to conduct the analyses described herein is available at <https://github.com/lockEF/bidifac>.

Acknowledgement

The authors gratefully acknowledges the support of NCI grant R21 CA231214.

A Algorithmic details

A.1 Fixed modules

For modules with fixed row and column sets defined by $\mathbf{R} : I \times K$ and $\mathbf{C} : J \times K$, the iterative estimation algorithm proceeds as follows:

1. Initialize $\hat{\mathbf{S}}_{\bullet\bullet}^{(k)} = \mathbf{0}_{M \times N}$ for $k = 1, \dots, K$.
2. For $k = 1, \dots, K$:
 - (a) Compute the residual matrix $\mathbf{X}_{\bullet\bullet}^{(k)} = \mathbf{X}_{\bullet\bullet} - \sum_{k' \neq k} \hat{\mathbf{S}}_{\bullet\bullet}^{(k')}$
 - (b) Set $\mathbf{X}_{ij}^{(k)} = \mathbf{0}_{M_i \times N_j}$ where $\mathbf{R}[i, k] = 0$ or $\mathbf{C}[j, k] = 0$
 - (c) Compute the SVD of $\mathbf{X}_{ij}^{(k)}$, $\mathbf{X}_{ij}^{(k)} = \mathbf{U}_{\cdot}^{(k)} \mathbf{D}^{(k)} \mathbf{V}_{\cdot}^{(k)}$
 - (d) Update $\hat{\mathbf{S}}_{\bullet\bullet}^{(k)} = \mathbf{U}_{\cdot}^{(k)} \hat{\mathbf{D}}^{(k)} \mathbf{V}_{\cdot}^{(k)}$ where $\hat{\mathbf{D}}[r, r] = \max(\mathbf{D}[r, r] - \lambda_k, 0)$ for $r = 1, 2, \dots$
3. Repeat step 2. until convergence of the objective function

$$\|\mathbf{X}_{\bullet\bullet} - \sum_{k=1}^K \hat{\mathbf{S}}_{\bullet\bullet}^{(k)}\|_F^2 + \sum_{k=1}^K 2 \left(\sqrt{\mathbf{M} \cdot \mathbf{R}[\cdot, k]} + \sqrt{\mathbf{N} \cdot \mathbf{C}[\cdot, k]} \right) \|\mathbf{S}_{\bullet\bullet}^{(k)}\|_*. \quad (18)$$

Step 2(d) finds minimizes the objective (18) for $\hat{\mathbf{S}}_{\bullet\bullet}^{(k)}$ given $\{\mathbf{S}_{\bullet\bullet}^{(k')}\}_{k' \neq k}$, by Proposition 2 in the main manuscript.

A.2 Undetermined modules

If the row and column sets defining the modules $\mathbf{R} : I \times K$ and $\mathbf{C} : J \times K$ are not predefined, we update them via a forward selection search process within the algorithm. The iterative estimation algorithm proceeds as follows:

1. Initialize $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}_{M \times N}$ for $k = 1, \dots, K$.
2. Initialize $\hat{\mathbf{C}}[j, k] = 1$ for $j = 1, \dots, J$.
3. For $k = 1, \dots, K$:
 - (a) Compute the residual matrix $\mathbf{X}_{\cdot\cdot}^{(k)} = \mathbf{X}_{\cdot\cdot} - \sum_{k' \neq k} \hat{\mathbf{S}}_{\cdot\cdot}^{(k')}$
 - (b) Update $\hat{\mathbf{R}}[\cdot, k]$ and $\hat{\mathbf{C}}[\cdot, k]$ as follows:
 - i. With $\hat{\mathbf{C}}[\cdot, k]$ fixed, update $\hat{\mathbf{R}}[\cdot, k]$ by forward selection, beginning with $\hat{\mathbf{R}}[\cdot, k] = \mathbf{0}$ and iteratively adding rows i ($\hat{\mathbf{R}}[i, k] = 1$) to minimize the objective (18).
 - ii. With $\hat{\mathbf{R}}[\cdot, k]$ fixed, update $\hat{\mathbf{C}}[\cdot, k]$ by forward selection, beginning with $\hat{\mathbf{R}}[\cdot, k] = \mathbf{0}$ and iteratively adding columns j ($\hat{\mathbf{C}}[j, k] = 1$) to minimize the objective (18).
 - iii. Repeat steps i. and ii. until convergence of the chosen row and column sets $\hat{\mathbf{C}}[\cdot, k]$ and $\hat{\mathbf{R}}[\cdot, k]$.
 - (c) Set $\mathbf{X}_{ij}^{(k)} = \mathbf{0}_{M_i \times N_j}$ where $\hat{\mathbf{R}}[i, k] = 0$ or $\hat{\mathbf{C}}[j, k] = 0$
 - (d) Compute the SVD of $\mathbf{X}_{ij}^{(k)}$, $\mathbf{X}_{ij}^{(k)} = \mathbf{U}_{\cdot}^{(k)} \mathbf{D}^{(k)} \mathbf{V}_{\cdot}^{(k)}$
 - (e) Update $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{U}_{\cdot}^{(k)} \hat{\mathbf{D}}^{(k)} \mathbf{V}_{\cdot}^{(k)}$ where $\hat{\mathbf{D}}[r, r] = \max(\mathbf{D}[r, r] - \lambda_k, 0)$ for $r = 1, 2, \dots$
4. Repeat step 3. until convergence of the objective function.

The forward selection steps in 3(b) can be performed relatively efficiently by only computing the singular values of the relevant submatrices of $\mathbf{X}_{\cdot\cdot}^{(k)}$, rather than the full SVD.

A.3 Tempered regularization

In practice, we find that the convergence of the algorithm in (A.2) improves substantially if the initial iterations use a high nuclear norm penalty that gradually decreases to the desired level of penalization. Thus, in our implementation for the first iteration the penalties are set to $\tilde{\lambda}_k = \alpha \lambda_k$ for $k = 1, \dots, K$ and some $\alpha > 1$. The penalties then gradually decrease over each subsequent iteration of the algorithm, before reaching the desired level of regularization ($\alpha = 1$).

B Module enumeration

As the default representation of model (6) in the main manuscript, set $K = (2^I - 1)(2^J - 1)$ and let \mathbf{R} and \mathbf{C} enumerate all possible modules as follows. For $k = 1, \dots, K$, let $\mathbf{R}[\cdot, k]$ be the I -digit binary representation for $k \bmod (2^I - 1) + 1$, where \bmod gives the modulo (remainder) operator. For $k = 1, \dots, K$, let $\mathbf{C}[\cdot, k]$ give the J -digit binary representation for $\lceil k / (2^I - 1) \rceil$, where $\lceil \cdot \rceil$ gives the ceiling operator.

C Proofs

Proposition 8. *Under objective (7) in the main manuscript, the following are necessary to allow for each $\hat{\mathbf{S}}_{\cdot}^{(k)}$ to be non-zero*

1. *If for $k' \neq k$ the rows and columns of module k' are contained within those for module k , $\mathbf{R}[i, k] - \mathbf{R}[i, k'] \geq 0 \forall i$ and $\mathbf{C}[j, k] - \mathbf{C}[j, k'] \geq 0 \forall j$, then $\lambda_k > \lambda_{k'}$.*
2. *If $\mathcal{I}_k \subset \{1, \dots, k - 1, k + 1, \dots, K\}$ is any subset of modules that together cover the rows and columns of module k , $\sum_{j \in \mathcal{I}_k} \mathbf{R}[\cdot, j] = r \cdot \mathbf{R}[\cdot, k]$ and $\sum_{j \in \mathcal{I}_k} \mathbf{C}[\cdot, j] = c \cdot \mathbf{C}[\cdot, k]$*

for positive integers r and c , then $\lambda_k < \sum_{j \in \mathcal{I}_k} \lambda_j$.

Proof. Let $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\mathbf{X}}$ be a minimizer of the objective function $f(\bullet)$. Assume a violation of condition 1., wherein $\lambda_{k'} \geq \lambda_k$. Consider another minimizer $\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$, where $\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}$ and $\tilde{\mathbf{S}}_{\cdot\cdot}^{(k')} = \hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + \hat{\mathbf{S}}_{\cdot\cdot}^{(k')}$, and all other modules are equal. Then, using the triangle inequality,

$$\begin{aligned} f(\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) - f(\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) &= \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \lambda_{k'} \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k')}\|_* - \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + \hat{\mathbf{S}}_{\cdot\cdot}^{(k')}\|_* \\ &\geq \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \lambda_{k'} \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k')}\|_* - \lambda_k (\|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k')}\|_*) \\ &\geq \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \lambda_{k'} \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k')}\|_* - \lambda_k (\|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_*) - \lambda_{k'} \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k')}\|_* \\ &= 0, \end{aligned}$$

and thus there is a solution in which module k is $\mathbf{0}$, regardless of the data $\mathbf{X}_{\cdot\cdot}$.

Now assume a violation of condition 2., wherein $\lambda_k \geq \sum_{j \in \mathcal{I}_k} \lambda_j$. Let $\hat{\mathbf{S}}^{(k)} = \sum_{j \in \mathcal{I}_k} \hat{\mathbf{S}}_{\cdot\cdot}^{j'}$, where $\hat{\mathbf{S}}_{\cdot\cdot}^{j'}$ contains the submatrix of $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}$ corresponding to $\mathbf{R}[\cdot, j]$ and $\mathbf{C}[\cdot, j]$ and $\mathbf{0}$ otherwise. Consider another decomposition $\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$, where $\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}$ and $\tilde{\mathbf{S}}_{\cdot\cdot}^{(j)} = \hat{\mathbf{S}}_{\cdot\cdot}^{(j)} + \hat{\mathbf{S}}_{\cdot\cdot}^{(j)'}$ for all $j \in \mathcal{I}_k$. Then,

$$\begin{aligned} f(\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) - f(\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) &= \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(j)}\|_* - \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(j)} + \hat{\mathbf{S}}_{\cdot\cdot}^{(j)'}\|_* \\ &\geq \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(j)}\|_* - \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(j)}\|_* - \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(j)'}\|_* \\ &= \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* - \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(j)'}\|_* \\ &\geq \lambda_k \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* - \sum_{j \in \mathcal{I}_k} \lambda_j \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* \\ &\geq 0, \end{aligned}$$

and thus there is a solution in which module k is $\mathbf{0}$, regardless of the data $\mathbf{X}_{\cdot\cdot}$. \square

Proposition 4. *Setting $\lambda_k = \sqrt{\mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, k] \cdot \mathbf{N}}$ satisfies the necessary conditions of Proposition 8.*

Proof. For condition 1., note that $\sqrt{\mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, k] \cdot \mathbf{N}} > \sqrt{\mathbf{R}[\cdot, j] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, j] \cdot \mathbf{N}}$.

For condition 2., note that

$$\begin{aligned} \sum_{j \in \mathcal{I}_k} \sqrt{\mathbf{R}[\cdot, j] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, j] \cdot \mathbf{N}} &\geq \sqrt{\sum_{j \in \mathcal{I}_k} \mathbf{R}[\cdot, j] \cdot \mathbf{M}} + \sqrt{\sum_{j \in \mathcal{I}_k} \mathbf{C}[\cdot, j] \cdot \mathbf{N}} \\ &= \sqrt{r \cdot \mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{c \cdot \mathbf{C}[\cdot, k] \cdot \mathbf{N}} \\ &> \sqrt{\mathbf{R}[\cdot, k] \cdot \mathbf{M}} + \sqrt{\mathbf{C}[\cdot, k] \cdot \mathbf{N}} \end{aligned}$$

□

Lemmas 1 and 2 below are used to establish Proposition 5 of the main manuscript.

Lemma 1. *Take two decompositions $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and $\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\tilde{\mathbf{X}}}$, and assume that both minimize the structured nuclear norm penalty:*

$$f_{\text{pen}}(\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) = f_{\text{pen}}(\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) = \min_{\mathbb{S}_{\hat{\mathbf{X}}}} f_{\text{pen}}(\{\mathbf{S}_{\cdot\cdot}^{(k)}\}_{k=1}^K).$$

Then, for any $\alpha \in [0, 1]$,

$$\|\alpha \hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + (1 - \alpha) \tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* = \alpha \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + (1 - \alpha) \|\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_*$$

for $k = 1, \dots, K$.

Proof. Because $\mathbb{S}_{\hat{\mathbf{X}}}$ is a convex space and f_{pen} is a convex function, the set of minimizers of f_{pen} over $\mathbb{S}_{\hat{\mathbf{X}}}$ is also convex. Thus,

$$f_{\text{pen}}(\{\alpha \hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + (1 - \alpha) \tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) = \min_{\mathbb{S}_{\hat{\mathbf{X}}}} f_{\text{pen}}(\{\mathbf{S}_{\cdot\cdot}^{(k)}\}_{k=1}^K).$$

The result follows from the convex property of the nuclear norm operator, which implies that for any two matrices of equal size $\hat{\mathbf{A}}$ and $\tilde{\mathbf{A}}$,

$$\|\alpha\hat{\mathbf{A}} + (1 - \alpha)\tilde{\mathbf{A}}\|_* \leq \alpha\|\hat{\mathbf{A}}\|_* + (1 - \alpha)\|\tilde{\mathbf{A}}\|_*. \quad (19)$$

Applying (19) to each additive term in f_{pen} gives

$$\begin{aligned} f_{\text{pen}}\left(\{\alpha\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + (1 - \alpha)\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K\right) &\leq \alpha f_{\text{pen}}(\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) + (1 - \alpha)f_{\text{pen}}(\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) \\ &= \min_{\mathbb{S}_{\hat{\mathbf{X}}}} f_{\text{pen}}(\{\mathbf{S}_{\cdot\cdot}^{(k)}\}_{k=1}^K). \end{aligned} \quad (20)$$

Because $\{\alpha\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + (1 - \alpha)\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$, the inequality in (20) must be an equality, and it follows that the inequality (19) must be an equality for each penalized term in the decomposition. \square

Lemma 2. *Take two matrices $\hat{\mathbf{A}}$ and $\tilde{\mathbf{A}}$. If $\|\hat{\mathbf{A}} + \tilde{\mathbf{A}}\| = \|\hat{\mathbf{A}}\|_* + \|\tilde{\mathbf{A}}\|_*$, and $\mathbf{U}\mathbf{D}_+\mathbf{V}^T$ is the SVD of $\hat{\mathbf{A}} + \tilde{\mathbf{A}}$, then $\hat{\mathbf{A}} = \hat{\mathbf{U}}\hat{\mathbf{D}}\hat{\mathbf{V}}^T$ where $\hat{\mathbf{D}}$ is diagonal and $\|\hat{\mathbf{A}}\|_* = \|\hat{\mathbf{D}}\|_*$, and $\tilde{\mathbf{A}} = \mathbf{U}\tilde{\mathbf{D}}\mathbf{V}^T$ where $\tilde{\mathbf{D}}$ is diagonal and $\|\tilde{\mathbf{A}}\|_* = \|\tilde{\mathbf{D}}\|_*$.*

Proof. Here we use the fact that the spectral norm is dual to the nuclear norm (Fazel et al., 2001). That is, if $\sigma_1(\mathbf{Z})$ is the maximum singular value of \mathbf{Z} (i.e., the spectral norm), then

$$\|\mathbf{A}\|_* = \sup_{\sigma_1(\mathbf{Z})=1} \langle \mathbf{Z}, \mathbf{A} \rangle.$$

Thus,

$$\sup_{\sigma_1(\mathbf{Z})=1} \langle \mathbf{Z}, \hat{\mathbf{A}} + \tilde{\mathbf{A}} \rangle = \sup_{\sigma_1(\mathbf{Z})=1} \langle \mathbf{Z}, \hat{\mathbf{A}} \rangle + \sup_{\sigma_1(\mathbf{Z})=1} \langle \mathbf{Z}, \tilde{\mathbf{A}} \rangle. \quad (21)$$

By the properties of the SVD,

$$\langle \mathbf{U}\mathbf{V}^T, \tilde{\mathbf{A}} \rangle + \langle \mathbf{U}\mathbf{V}^T, \hat{\mathbf{A}} \rangle = \langle \mathbf{U}\mathbf{V}^T, \hat{\mathbf{A}} + \tilde{\mathbf{A}} \rangle = \sup_{\sigma_1(\mathbf{Z})=1} \langle \mathbf{Z}, \hat{\mathbf{A}} + \tilde{\mathbf{A}} \rangle. \quad (22)$$

By (21) and (22),

$$\langle \mathbf{UV}^T, \hat{\mathbf{A}} \rangle = \sup_{\sigma_1(\mathbf{Z})=1} \langle \mathbf{Z}, \mathbf{A} \rangle = \|\mathbf{A}\|_*,$$

and similarly $\langle \mathbf{UV}^T, \tilde{\mathbf{A}} \rangle = \|\tilde{\mathbf{A}}\|_*$. Let $\tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T$ give the SVD of $\tilde{\mathbf{A}}$. Note that

$$\langle \mathbf{UV}^T, \tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T \rangle = \text{Tr}(\mathbf{V}\mathbf{U}^T\tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T) = \text{Tr}(\mathbf{V}^T\tilde{\mathbf{V}}\mathbf{U}^T\tilde{\mathbf{U}}\tilde{\mathbf{D}}),$$

and

$$\text{Tr}(\mathbf{V}^T\tilde{\mathbf{V}}\mathbf{U}^T\tilde{\mathbf{U}}\tilde{\mathbf{D}}) = \|\tilde{\mathbf{A}}\|_* = \sum_i \tilde{\mathbf{D}}[i, i]$$

if and only if $\mathbf{V}^T\tilde{\mathbf{V}}\mathbf{U}^T\tilde{\mathbf{U}}[i, i] = 1$ where $\tilde{\mathbf{D}}[i, i] > 0$. It follows that the left and right singular vectors of $\tilde{\mathbf{A}}$ that correspond to non-zero singular values must also be singular vectors of $\hat{\mathbf{A}} + \tilde{\mathbf{A}}$. By an analogous argument, the left and right singular vectors that correspond to non-zero singular values in $\hat{\mathbf{A}}$ must also be singular vectors of $\hat{\mathbf{A}} + \tilde{\mathbf{A}}$. \square

Proposition 5. *Take two decompositions $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and $\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\tilde{\mathbf{X}}}$, and assume that both minimize the structured nuclear norm penalty:*

$$f_{pen}(\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) = f_{pen}(\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) = \min_{\mathbb{S}_{\hat{\mathbf{X}}}} f_{pen}(\{\mathbf{S}_{\cdot\cdot}^{(k)}\}_{k=1}^K).$$

Then, $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{U}_{\cdot}^{(k)}\hat{\mathbf{D}}^{(k)}\mathbf{V}_{\cdot}^{(k)T}$ and $\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{U}_{\cdot}^{(k)}\tilde{\mathbf{D}}^{(k)}\mathbf{V}_{\cdot}^{(k)T}$ where $\mathbf{U}_{\cdot}^{(k)} : M \times R_k$ and $\mathbf{V}_{\cdot}^{(k)} : N \times R_k$ have orthonormal columns, and $\hat{\mathbf{D}}^{(k)}$ and $\tilde{\mathbf{D}}^{(k)}$ are diagonal.

Proof. This result is a direct corollary of Lemmas 1 and 2. Lemma 1 implies $\|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + \tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* = \|\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_* + \|\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_*$ for each k , and then lemma 2 implies the result. \square

Theorem 2. *Consider $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and let $\mathbf{U}_{\cdot}^{(k)}\hat{\mathbf{D}}\mathbf{V}_{\cdot}^{(k)T}$ give the SVD of $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}$ for $k = 1, \dots, K$. The following three properties uniquely identify $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$.*

1. $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$ minimizes $f_{pen}(\cdot)$ over $\mathbb{S}_{\hat{\mathbf{X}}}$,

2. $\{\hat{\mathbf{U}}_i^{(k)}[\bullet, r] : \mathbf{R}[i, k] = 1 \text{ and } \hat{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $i = 1, \dots, I$,
3. $\{\hat{\mathbf{V}}_j^{(k)}[\bullet, r] : \mathbf{C}[j, k] = 1 \text{ and } \hat{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $j = 1, \dots, J$.

Proof. Take two decomposition $\{\hat{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K$ and $\{\tilde{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K$ that satisfy properties 1., 2., and 3.; we will show that $\{\hat{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K = \{\tilde{\mathbf{S}}_{\bullet}^{(k)}\}_{k=1}^K$. For each $k = 1, \dots, K$, write $\hat{\mathbf{S}}_{\bullet}^{(k)} = \mathbf{U}_{\bullet}^{(k)} \hat{\mathbf{D}} \mathbf{V}_{\bullet}^{(k)T}$ and $\tilde{\mathbf{S}}_{\bullet}^{(k)} = \mathbf{U}_{\bullet}^{(k)} \tilde{\mathbf{D}}^{(k)} \mathbf{V}_{\bullet}^{(k)T}$ as in Proposition 5. Then, it suffices to show that $\hat{\mathbf{D}}^{(k)}[r, r] = \tilde{\mathbf{D}}^{(k)}[r, r]$ for all k, r .

First, consider module $k = 1$ with $\mathbf{R}[\bullet, 1] = [1 \ 0 \ \dots \ 0]^T$ and $\mathbf{C}[\bullet, 1] = [1 \ 0 \ \dots \ 0]^T$. By way of contradiction, assume $\hat{\mathbf{D}}^{(1)}[1, 1] > 0$ and $\tilde{\mathbf{D}}^{(1)}[1, 1] = 0$. The linear independence of $\{\mathbf{V}_j^{(k)}[\bullet, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0\}$ and $\{\mathbf{V}_j^{(k)}[\bullet, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$ implies that

$$\text{row}(\mathbf{X}_{\bullet}) = \text{span}\{\mathbf{U}_{\bullet}^{(k)}[\bullet, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0\} = \text{span}\{\{\mathbf{U}_{\bullet}^{(k)}[\bullet, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}.$$

Thus, $\mathbf{U}^{(1)}[\bullet, 1] \in \text{span}\{\{\mathbf{U}_{\bullet}^{(k)}[\bullet, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$, and it follows from the orthogonality of $\mathbf{U}^{(1)}[\bullet, 1]$ and $\{\mathbf{U}^{(1)}[\bullet, r], r > 1\}$ that

$$\mathbf{U}^{(1)}[\bullet, 1] \in \text{span}\{\{\mathbf{U}_{\bullet}^{(k)}[\bullet, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0 \text{ and } k > 1\}.$$

Moreover, because $\mathbf{U}_i^{(1)} = \mathbf{0}$ for any $i > 1$ and $\{\mathbf{U}_i^{(k)}[\bullet, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent it follows that

$$\mathbf{U}^{(1)}[\bullet, 1] \in \text{span}\{\{\mathbf{U}_{\bullet}^{(k)}[\bullet, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0, k > 1, \text{ and } \mathbf{R}[i, k] = 0 \text{ for any } i > 1\}. \quad (23)$$

Note that (23) implies $\mathbf{U}_1^{(1)}[\bullet, 1] \in \text{row}(\mathbf{X}_{12} + \dots + \text{row}(\mathbf{X}_{1J}))$, however, this is contradicted by the linear independence of $\mathbf{U}_1^{(1)}[\bullet, 1]$ and $\{\mathbf{U}_i^{(k)}[\bullet, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0, k > 1\}$. Thus, we conclude that $\hat{\mathbf{D}}^{(1)}[1, 1] > 0$ implies $\tilde{\mathbf{D}}^{(1)}[1, 1] > 0$. Analogous arguments show that $\tilde{\mathbf{D}}^{(k)}[r, r] > 0$ if and only if $\hat{\mathbf{D}}^{(k)}[r, r] > 0$ for any pair (r, k) . It follows that $\{\mathbf{U}_i^{(k)}[\bullet, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0 \text{ or } \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $i = 1, \dots, I$, and $\{\mathbf{V}_j^{(k)}[\bullet, r] :$

$\hat{\mathbf{D}}^{(k)}[r, r] > 0$ or $\tilde{\mathbf{D}}^{(k)}[r, r] > 0$ are linearly independent for $j = 1, \dots, J$. Thus,

$$\sum_{k=1}^K \mathbf{U}^{(k)} (\hat{\mathbf{D}}^{(k)} - \tilde{\mathbf{D}}^{(k)}) \mathbf{V}^{(k)T} = \sum_{k=1}^K \hat{\mathbf{S}}_{..}^{(k)} - \tilde{\mathbf{S}}_{..}^{(k)} = \mathbf{X}_{..} - \mathbf{X}_{..} = \mathbf{0}$$

implies that $\hat{\mathbf{D}}^{(k)}[r, r] = \tilde{\mathbf{D}}^{(k)}[r, r]$ for all k, r .

□

References

- Akbani, R., Ng, P. K. S., Werner, H. M., Shahmoradgoli, M., Zhang, F., Ju, Z., Liu, W., Yang, J.-Y., Yoshihara, K., Li, J., et al. (2014). A pan-cancer proteomic perspective on the cancer genome atlas. *Nature communications*, 5:3887.
- 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(6).
- Fazel, M., Hindi, H., Boyd, S. P., et al. (2001). A rank minimization heuristic with application to minimum order system approximation. In *Proceedings of the American control conference*, volume 6, pages 4734–4739. Citeseer.
- Gabasova, E., Reid, J., and Wernisch, L. (2017). Clusternomics: Integrative context-dependent clustering for heterogeneous datasets. *PLoS computational biology*, 13(10):e1005781.
- Gavish, M. and Donoho, D. L. (2017). Optimal shrinkage of singular values. *IEEE Transactions on Information Theory*, 63(4):2137–2152.

- Gaynanova, I. and Li, G. (2019). Structural learning and integrative decomposition of multi-view data. *Biometrics*.
- Hellton, K. H. and Thoresen, M. (2016). Integrative clustering of high-dimensional data with joint and individual clusters. *Biostatistics*, 17(3):537–548.
- Hoadley, K. A., Yau, C., Hinoue, T., Wolf, D. M., Lazar, A. J., Drill, E., Shen, R., Taylor, A. M., Cherniack, A. D., Thorsson, V., et al. (2018). Cell-of-origin patterns dominate the molecular classification of 10,000 tumors from 33 types of cancer. *Cell*, 173(2):291–304.
- Hoadley, K. A., Yau, C., Wolf, D. M., Cherniack, A. D., Tamborero, D., Ng, S., Leiserson, M. D., Niu, B., McLellan, M. D., Uzunangelov, V., et al. (2014). Multiplatform analysis of 12 cancer types reveals molecular classification within and across tissues of origin. *Cell*, 158(4):929–944.
- Huo, Z. and Tseng, G. (2017). Integrative sparse k-means with overlapping group lasso in genomic applications for disease subtype discovery. *The annals of applied statistics*, 11(2):1011.
- Hutter, C. and Zenklusen, J. C. (2018). The Cancer Genome Atlas: creating lasting value beyond its data. *Cell*, 173(2):283–285.
- Kandoth, C., McLellan, M. D., Vandin, F., Ye, K., Niu, B., Lu, C., Xie, M., Zhang, Q., McMichael, J. F., Wyczalkowski, M. A., et al. (2013). Mutational landscape and significance across 12 major cancer types. *Nature*, 502(7471):333.
- Kurucz, M., Benczúr, A. A., and Csalogány, K. (2007). Methods for large scale svd with missing values. In *Proceedings of KDD cup and workshop*, volume 12, pages 31–38.

- Li, G., Gaynanova, I., et al. (2018). A general framework for association analysis of heterogeneous data. *The Annals of Applied Statistics*, 12(3):1700–1726.
- Li, G. and Jung, S. (2017). Incorporating covariates into integrated factor analysis of multi-view data. *Biometrics*, 73(4):1433–1442.
- Lock, E. F. and Dunson, D. B. (2013). Bayesian consensus clustering. *Bioinformatics*, 29(20):2610–2616.
- Lock, E. F., Hoadley, K. A., Marron, J., and Nobel, A. B. (2013). Joint and Individual Variation Explained (JIVE) for integrated analysis of multiple data types. *The Annals of Applied Statistics*, 7(1):523.
- Mazumder, R., Hastie, T., and Tibshirani, R. (2010). Spectral regularization algorithms for learning large incomplete matrices. *Journal of machine learning research*, 11(Aug):2287–2322.
- Mo, Q., Shen, R., Guo, C., Vannucci, M., Chan, K. S., and Hilsenbeck, S. G. (2017). A fully bayesian latent variable model for integrative clustering analysis of multi-type omics data. *Biostatistics*, 19(1):71–86.
- Noushmehr, H., Weisenberger, D. J., Diefes, K., Phillips, H. S., Pujara, K., Berman, B. P., Pan, F., Pelloski, C. E., Sulman, E. P., Bhat, K. P., et al. (2010). Identification of a cpg island methylator phenotype that defines a distinct subgroup of glioma. *Cancer cell*, 17(5):510–522.
- O’Connell, M. and Lock, E. F. (2017). Linked matrix factorization. *arXiv preprint arXiv:1707.06573*.

- O’Connell, M. J. and Lock, E. F. (2016). R. jive for exploration of multi-source molecular data. *Bioinformatics*, 32(18):2877–2879.
- Park, J. Y. and Lock, E. F. (2019). Integrative factorization of bidimensionally linked matrices. *Biometrics*.
- Rudelson, M. and Vershynin, R. (2010). Non-asymptotic theory of random matrices: extreme singular values. In *Proceedings of the International Congress of Mathematicians 2010 (ICM 2010) (In 4 Volumes) Vol. I: Plenary Lectures and Ceremonies Vols. II–IV: Invited Lectures*, pages 1576–1602. World Scientific.
- Shabalin, A. A. and Nobel, A. B. (2013). Reconstruction of a low-rank matrix in the presence of gaussian noise. *Journal of Multivariate Analysis*, 118:67–76.
- Shen, R., Wang, S., and Mo, Q. (2013). Sparse integrative clustering of multiple omics data sets. *The annals of applied statistics*, 7(1):269.
- TCGA Research Network (2015). Comprehensive, integrative genomic analysis of diffuse lower-grade gliomas. *New England Journal of Medicine*, 372(26):2481–2498.
- TCGA Research Network et al. (2012). Comprehensive molecular portraits of human breast tumors. *Nature*, 490(7418):61.
- TCGA Research Network et al. (2014). Comprehensive molecular profiling of lung adenocarcinoma. *Nature*, 511(7511):543.
- Verhaak, R. G., Hoadley, K. A., Purdom, E., Wang, V., Qi, Y., Wilkerson, M. D., Miller, C. R., Ding, L., Golub, T., Mesirov, J. P., et al. (2010). Integrated genomic analysis identifies clinically relevant subtypes of glioblastoma characterized by abnormalities in PDGFRA, IDH1, EGFR, and NF1. *Cancer cell*, 17(1):98–110.

- Weinstein, J. N., Collisson, E. A., Mills, G. B., Shaw, K. R. M., Ozenberger, B. A., Ellrott, K., Shmulevich, I., Sander, C., Stuart, J. M., Network, C. G. A. R., et al. (2013). The cancer genome atlas pan-cancer analysis project. *Nature genetics*, 45(10):1113–1120.
- Yang, Z. and Michailidis, G. (2016). A non-negative matrix factorization method for detecting modules in heterogeneous omics multi-modal data. *Bioinformatics*, 32(1):1–8.
- Zack, T. I., Schumacher, S. E., Carter, S. L., Cherniack, A. D., Saksena, G., Tabak, B., Lawrence, M. S., Zhang, C.-Z., Wala, J., Mermel, C. H., et al. (2013). Pan-cancer patterns of somatic copy number alteration. *Nature genetics*, 45(10):1134–1140.
- Zhu, H., Li, G., and Lock, E. F. (2019). Generalized integrative principal component analysis for multi-type data with block-wise missing structure. *Biostatistics*.