

Empirical Bayes Linked Matrix Decomposition

Eric F. Lock^{1*}

^{1*}Division of Biostatistics and Health Data Science, School of Public Health, University of Minnesota, Minneapolis, 55455, MN, USA.

Corresponding author(s). E-mail(s): elock@umn.edu;

Abstract

Data for several applications in diverse fields can be represented as multiple matrices that are linked across rows or columns. This is particularly common in molecular biomedical research, in which multiple molecular “omics” technologies may capture different feature sets (e.g., corresponding to rows in a matrix) and/or different sample populations (corresponding to columns). This has motivated a large body of work on integrative matrix factorization approaches that identify and decompose low-dimensional signal that is shared across multiple matrices or specific to a given matrix. We propose an empirical variational Bayesian approach to this problem that has several advantages over existing techniques, including the flexibility to accommodate shared signal over any number of row or column sets (i.e., bidimensional integration), an intuitive model-based objective function that yields appropriate shrinkage for the inferred signals, and a relatively efficient estimation algorithm with no tuning parameters. A general result establishes conditions for the uniqueness of the underlying decomposition for a broad family of methods that includes the proposed approach. For scenarios with missing data, we describe an associated iterative imputation approach that is novel for the single-matrix context and a powerful approach for “blockwise” imputation (in which an entire row or column is missing) in various linked matrix contexts. Extensive simulations show that the method performs very well under different scenarios with respect to recovering underlying low-rank signal, accurately decomposing shared and specific signals, and accurately imputing missing data. The approach is applied to gene expression and miRNA data from breast cancer tissue and normal breast tissue, for which it gives an informative decomposition of variation and outperforms alternative strategies for missing data imputation.

Keywords: Data integration, dimension reduction, low-rank factorization, missing data imputation, variational Bayes

1 Introduction

Low-rank matrix factorization techniques are widely used for many machine learning tasks such as data compression and dimension reduction, denoising to approximate underlying signal, and matrix completion. Often, instead of a single matrix, the data for a given application takes the form of multiple matrices that are linked by rows and/or columns. For example, in Section 9 we consider data from The Cancer Genome Atlas, in which gene expression (mRNA) and microRNA (miRNA) data are available from different platforms for breast cancer tumor samples and normal breast tissue samples from unrelated individuals. Thus, the data can be represented as four matrices: (1) mRNA for cancer tissues, (2) mRNA for normal tissues, (3) miRNA for cancer tissues, and (4) miRNA for normal tissues; these matrices share both rows (here, molecular features) columns (here, samples), giving a *bidimensionally linked* structure. Furthermore, these data have *column-wise* missingness, in which molecular data for one of the platforms (mRNA or miRNA) is entirely missing for some samples in each of the cancer and control cohorts. We are interested in imputing the missing data, as well as identifying patterns of variation that are shared or specific to the different data sources and sample cohorts; for these tasks, a low-rank factorization approach that integrates all of the matrices together while accounting for their differences is attractive.

There is a growing and dynamic literature on multi-matrix low-rank factorization approaches. The majority of existing methods assume that just one dimension, rows or columns, is shared across matrices, i.e., *unidimensional* integration. For example, the Joint and Individual Variation Explained (JIVE) approach (Lock et al., 2013; O’Connell and Lock, 2016) decomposes low-rank structure that is shared across matrices from variation that is specific to each matrix. Several related approaches yield a similar decomposition using different strategies or assumptions (Feng et al., 2018; Schouteden et al., 2014; Yang and Michailidis, 2016; Tang and Allen, 2021), while other approaches such as structural learning and integrative decomposition (SLIDE) (Gaynanova and Li, 2019) allow for structure that is partially shared across any subset of matrices. Such an approach is efficient and useful for interpretation, as for many applications identifying shared, partially shared, and specific underlying structure is of interest.

The bidimensional linked matrix factorization (BIDIFAC) approach (Park and Lock, 2020; Lock et al., 2022) was developed to accommodate bidimensionally linked integration scenarios, by identifying structure that is shared or specific across both row sets and column sets. The objective function for BIDIFAC involves a structured nuclear norm penalty, which approximates the data as a decomposition of low-rank *modules* that are shared across row or column subsets. This approach is attractive in the bidimensionally linked context because it has favorable convergence properties and generally leads to a uniquely-defined decomposition (Lock et al., 2022); methods for unidimensional integration without shrinkage penalties typically require orthogonality (Lock et al., 2013; Feng et al., 2018; Gaynanova and Li, 2019; Yi et al., 2023) or other constraints to assure uniqueness, which do not extend to the bidimensionally linked context. Furthermore, the application of the structured nuclear norm penalty for

unidimensional integration (termed UNIFAC) has been shown to outperform alternative approaches with respect to accurately decomposing shared and specific structures (Lock et al., 2022) and imputing missing data (Park and Lock, 2020) under certain scenarios. However, a drawback is that it tends to over-penalize the estimated structure, and is thus biased toward an over-conservative solution with respect to the underlying low-rank structure (Yi et al., 2023).

For a single matrix the tendency of the nuclear norm to over-shrink low-rank signal has been well-documented, and in contrast low-rank approximations without any further shrinkage tend to overfit (Josse and Sardy, 2016). Empirical Bayes approaches to low-rank factorization are an attractive compromise, in which the optimal levels of shrinkage are determined empirically within a model-based framework (Wang and Stephens, 2021). Under a variational Bayes approach, the true posterior distribution for Bayesian model is approximated by minimizing the free energy (reducing to Kullback-Leibler divergence) under simpler distributional assumptions (Fox and Roberts, 2012). In particular, an empirical variational Bayes approach to low-rank matrix factorization under an assumption of normally distributed factor matrices has desirable theoretical properties, and the resulting approximation for a scenario without missing data can be efficiently computed via a closed-form shrinkage formula on the observed singular values of the matrix (Nakajima et al., 2012, 2013).

In this article, we describe an empirical variational Bayesian approach to identify and decompose shared, partially shared, or specific low-rank structure in the context of unidimensionally or bidimensionally linked matrices. Moreover, we describe an accompanying iterative imputation approach analogous to an expectation-maximization (EM) algorithm that can handle entry-wise missing or row/column-wise missing values. The proposed approach is free of any tuning parameters, and is shown to share desirable properties of the BIDIFAC approach while addressing the issue of over-shrinkage.

2 Notation and Setting

Throughout this article bold uppercase characters (\mathbf{X}) correspond to matrices and bold lowercase characters (\mathbf{x}) correspond to vectors. Square brackets index the entries within a vector or matrix, e.g., $\mathbf{X}[m, n]$. The singular value decomposition (SVD) of a matrix $\mathbf{X} : M \times N$ is given by $\mathbf{U}_\mathbf{X} \mathbf{D}_\mathbf{X} \mathbf{V}_\mathbf{X}^T$, where the diagonal entries of $\mathbf{D}_\mathbf{X}$, $\mathbf{D}_\mathbf{X}[r, r]$, are the singular values of \mathbf{X} ordered from largest to smallest for $r = 1, \dots, \min(M, N)$. The Frobenius norm is given by $\|\cdot\|_F$, so that $\|\mathbf{X}\|_F^2$ is the sum of the squared entries in \mathbf{X} ; the nuclear norm is given by $\|\cdot\|_*$, which is the sum of the singular values in a matrix:

$$\|\mathbf{X}\|_* = \sum_{r=1}^{\min(M, N)} \mathbf{D}_\mathbf{X}[r, r].$$

We consider the setting of a collection of matrices $\mathbf{X}_{ij} : M_i \times N_j$ for $i = 1, \dots, I$ and $j = 1, \dots, J$. The column sets N_1, \dots, N_J and row sets M_1, \dots, M_L are consistent across the matrices. The subscript ‘ \cdot ’ defines concatenation across row or column sets, and the full data \mathbf{X}_\cdot can be represented as a single $M \times N$ matrix where $M =$

$M_1 + \dots + M_I$ and $N = N_1 + \dots + N_J$:

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

As for our motivating application in Section 9, the I column sets may correspond to different sample cohorts and the J row sets may correspond to features from different high-dimensional sources.

3 Single-matrix results

Here we review some established results for a single matrix that are critical to what follows. Propositions 1 and 2 describe solutions to the least squares matrix approximation problem under a rank-constrained or nuclear norm penalized criterion, respectively.

Proposition 1. For $\mathbf{X} : M \times N$, the minimizer of the least squares objective $\|\mathbf{X} - \mathbf{S}\|_F^2$ under the constraint $\text{rank}(\mathbf{S}) = R \leq \min(M, N)$ is given by $\mathbf{S} = \mathbf{U}_\mathbf{X} \mathbf{D}_\mathbf{S} \mathbf{V}_\mathbf{X}^T$ where $\mathbf{D}_\mathbf{S}$ is diagonal with $\mathbf{D}_\mathbf{S}[r, r] = \mathbf{D}_\mathbf{X}[r, r]$ for $r = 1, \dots, R$ and $\mathbf{D}_\mathbf{S}[r, r] = 0$ for $r > R$.

Proposition 2. For $\mathbf{X} : M \times N$, the minimizer of the nuclear norm penalized objective $\frac{1}{2} \|\mathbf{X} - \mathbf{S}\|_F^2 + \lambda \|\mathbf{S}\|_*$ is given by $\mathbf{S} = \mathbf{U}_\mathbf{X} \mathbf{D}_\mathbf{S} \mathbf{V}_\mathbf{X}^T$ where $\mathbf{D}_\mathbf{S}$ is diagonal with $\mathbf{D}_\mathbf{S}[r, r] = \max(\mathbf{D}_\mathbf{X}[r, r] - \lambda, 0)$ for all r .

Proposition 1 is well-known, and a proof of Proposition 2 can be found in Mazumder et al. (2010). Because of how they operate on the singular values \mathbf{X} , Proposition 1 represents a *hard-thresholding* and Proposition 2 a *soft-thresholding* approach to low-rank approximation. Proposition 3 below establishes that the soft-thresholding operator also solves an objective with L_2 -penalties on the factor components of \mathbf{S} , while Corollary 4 further establishes that it is the posterior mode of a Bayesian model with normal priors on the factor components.

Proposition 3. For $\mathbf{X} : M \times N$,

$$\min_{\mathbf{U}:M \times H, \mathbf{V}:N \times H} \|\mathbf{X} - \mathbf{U}\mathbf{V}^T\|_F^2 + \lambda(\|\mathbf{U}\|_F^2 + \|\mathbf{V}\|_F^2) = \min_{\mathbf{S}:\text{rank}(\mathbf{S}) \leq H} \|\mathbf{X} - \mathbf{S}\|_F^2 + 2\lambda\|\mathbf{S}\|_* \quad (2)$$

Moreover, $\widehat{\mathbf{S}} = \widehat{\mathbf{U}}\widehat{\mathbf{V}}^T$, where $\widehat{\mathbf{S}}$ solves the right-hand side of (2) and $\{\widehat{\mathbf{U}}, \widehat{\mathbf{V}}\}$ solves the left-hand side of (2).

Corollary 4. For $\mathbf{X} : M \times N$, if $\mathbf{U} : M \times H$ and $\mathbf{V} : N \times H$ have independent $\text{Normal}(0, \sigma^2/\lambda)$ entries, and $\mathbf{X} = \mathbf{U}\mathbf{V}^T + \mathbf{E}$ where \mathbf{E} has independent $\text{Normal}(0, \sigma^2)$ entries, then the mode of the posterior distribution $p(\mathbf{U}, \mathbf{V} \mid \mathbf{X})$ satisfies $\widehat{\mathbf{S}} = \widehat{\mathbf{U}}\widehat{\mathbf{V}}^T$, where $\widehat{\mathbf{S}}$ solves the right-hand side of (2).

A proof of Proposition 3 can be found in Mazumder et al. (2010), and Corollary 4 follows by noting that the log of the posterior density for the specified model is proportional to the left-hand side of (2). While the results hold for general H , in what follows we set H as the largest possible rank $H = \min(M, N)$; the actual estimated rank will tend to be smaller due to singular value thresholding (as in Proposition 2). The distribution of singular values for the error matrix \mathbf{E} has been well-studied, including the following results on the first singular value under general conditions from Rudelson and Vershynin (2010):

Proposition 5. *If $\mathbf{E} : M \times N$ has independent entries with mean 0, variance σ^2 and finite fourth moment, then $\mathbf{D}[1, 1] \approx \sigma(\sqrt{M} + \sqrt{N})$ as $M, N \rightarrow \infty$; further, if the entries are Gaussian, $\mathbb{E}(\mathbf{D}[1, 1]) \leq \sigma(\sqrt{M} + \sqrt{N})$ for any M, N .*

These results can motivate choice of the nuclear norm penalty λ . A conservative approach, if the error variance σ^2 is known, is to select $\lambda = \sigma(\sqrt{M} + \sqrt{N})$; per Propositions 5 and 2, this will tend to shrink to 0 any singular values in \mathbf{X} that result from error with no true low-rank signal. Another approach is to adopt an empirical Bayes estimate for λ based on the model in Corollary 4. However, a general limitation of the nuclear norm approach to low-rank matrix reconstruction is the use of a universal threshold on all singular values; in practice larger singular values should be penalized less, as they capture a relatively lower ratio of error and higher signal (Shabalin and Nobel, 2013). An analogous interpretation is that the entries of \mathbf{U} and \mathbf{V} should not have identical variance across columns, and this can be relaxed by a model in which they have column-specific variances: $\mathbf{U}[\cdot, r] \sim \text{Normal}(\mathbf{0}, \tau_{\mathbf{U}, r}^2 \mathbf{I})$ and $\mathbf{V}[\cdot, r] \sim \text{Normal}(\mathbf{0}, \tau_{\mathbf{V}, r}^2 \mathbf{I})$ for $r = 1, \dots, H$.¹

For given $\boldsymbol{\tau} = \{\tau_{\mathbf{U}, r} \tau_{\mathbf{V}, r}\}_{r=1}^R$, the approximation resulting from the posterior mode for $p(\mathbf{U}, \mathbf{V} \mid \mathbf{X}, \boldsymbol{\tau})$ is given by thresholding the singular values of \mathbf{X} , analogous to Corollary 4. To determine $\boldsymbol{\tau}$, directly maximizing the joint likelihood $p(\mathbf{X}, \mathbf{U}, \mathbf{V} \mid \boldsymbol{\tau})$ over \mathbf{U}, \mathbf{V} and $\boldsymbol{\tau}$ degenerates to $\boldsymbol{\tau} \rightarrow \mathbf{0}$ in the global solution, though it can have a reasonable local mode that is a singular value thresholding operator (Nakajima and Sugiyama, 2014). Further, a marginal empirical Bayes estimator that optimizes $p(\mathbf{X} \mid \boldsymbol{\tau})$ (effectively integrating over \mathbf{U} and \mathbf{V}) is difficult to obtain computationally or analytically. As an alternative, one can use an empirical variational Bayes (EVB) approach. In general, for a variational Bayes approach (Attias, 1999), the true posterior of a Bayesian model with parameters Θ and data \mathbf{y} , $p(\Theta \mid \mathbf{y})$ is approximated by another distribution $q(\Theta)$ under simplifying assumptions to minimize the free energy

$$F(q) = E_q \log \frac{q(\Theta)}{p(\Theta, \mathbf{y})} = E_q \log \frac{q(\Theta)}{p(\Theta \mid \mathbf{y})} - \log p(\mathbf{y}),$$

where $E_q(\cdot)$ is the expected value with respect to the distribution defined by $q(\cdot)$. In the context of low-rank approximation, a useful simplifying assumption for $q(\cdot)$ is that the left and right factor matrices are independent, $q(\mathbf{U}, \mathbf{V}) = q_u(\mathbf{U})q_v(\mathbf{V})$. Thus, for

¹For most scenarios $\tau_{\mathbf{U}, r}^2$ and $\tau_{\mathbf{V}, r}^2$ are not independently identifiable, as the level of signal is determined by their product.

hyperparameters $\boldsymbol{\tau}$ and σ the free energy to minimize is

$$F(q | \sigma, \boldsymbol{\tau}) = E_q \log \frac{q_u(\mathbf{U})q_v(\mathbf{V})}{p(\mathbf{X} | \mathbf{U}, \mathbf{V}, \sigma)p(\mathbf{U} | \boldsymbol{\tau}_U)p(\mathbf{V} | \boldsymbol{\tau}_V)}, \quad (3)$$

and a point estimate for low-rank structure $\hat{\mathbf{S}}$ can be obtained via its expected value under $q(\cdot)$:

$$\hat{\mathbf{S}} = E_q(\mathbf{U}\mathbf{V}^T). \quad (4)$$

Under an empirical variational Bayes approach, $\boldsymbol{\tau}$ and σ are also estimated by minimizing the free energy in (3). For fully-observed \mathbf{X} , Nakajima et al. (2012) showed that the global solution to the empirical variation Bayes objective can be obtained in closed form as a singular value thresholding operator; their result for fixed σ is given in Theorem 6, and for estimation of σ in Theorem 7.

Theorem 6. *The approximation $\hat{\mathbf{S}}$ (4) obtained by minimizing (3) over q and $\boldsymbol{\tau}$ is $\hat{\mathbf{S}} = \mathbf{S} = \mathbf{U}_X \mathbf{D}_S \mathbf{V}_X^T$ where \mathbf{D}_S is diagonal with*

$$\mathbf{D}_S[r, r] = 0 \text{ if } D_{\mathbf{X}}[r, r] < \sigma \sqrt{M + N + \sqrt{MN}(\kappa + 1/\kappa)}$$

and

$$\mathbf{D}_S[r, r] = \left(\mathbf{D}_{\mathbf{X}}[r, r]^2 - (M + N)\sigma^2 + \sqrt{(\mathbf{D}_{\mathbf{X}}[r, r]^2 - (M + N)\sigma^2)^2 - 4MN\sigma^4} \right) / (2\mathbf{D}_{\mathbf{X}}[r, r]) \text{ otherwise.}$$

The value κ is given as the zero-cross point of the function

$$f(x) = x\sqrt{N/M} \log \left(x\sqrt{M/N} + 1 \right) + x\sqrt{M/N} \log \left(x\sqrt{N/M} + 1 \right) - 1.$$

Theorem 7. *Assume without loss of generality that $M > N$ and $\alpha = N/M$. The estimate $\hat{\sigma}$ that minimizes (3) is the global minimizer of*

$$\Psi(\sigma) = \sum_{r=1}^N \psi_1 \left(\frac{\mathbf{D}_X[r, r]^2}{M\sigma^2} \right),$$

where

$$\begin{aligned} \psi_2(x) &= x - \log x + \mathbf{1}_{\{x > c\}} \psi_2(x), \\ \psi_2(x) &= \log(\sqrt{\alpha}\psi_3(x) + 1) + \alpha \log(\psi_3(x)/\sqrt{\alpha} + 1) - \alpha\psi_3(x), \end{aligned}$$

and

$$\psi_3(x) = \frac{1}{2\sqrt{\alpha}} \left(x - 1 - \alpha + \sqrt{(x - (1 + \alpha))^2 - 4\alpha} \right).$$

Together, Theorems 7 and 6 provide a powerful and efficient way to approximate low-rank signal without any tuning parameters. The only substantial computing

requirement is in obtaining the SVD of \mathbf{X} , as other steps involve univariate functions that are trivial to solve.

4 Bidimensional linked matrix factorization

BIDIFAC+ (Lock et al., 2022) was developed to decompose bidimensionally linked data as in (1) into modules of low-rank structure that are shared across column sets and/or row sets. That is,

$$\mathbf{X}_{..} \approx \sum_{k=1}^K \mathbf{S}_{..}^{(k)}, \quad (5)$$

where each $\mathbf{S}_{..}^{(k)}$ is a concatenation of blocks $\mathbf{S}_{ij}^{(k)}$ as in (1) and

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

Here, \mathbf{R} and \mathbf{C} are binary matrices that indicate the presence of module k ($\mathbf{S}_{..}^{(k)}$) across the row and column sets; that is, $\mathbf{S}_{..}^{(k)}$ is low-rank structure specific to the submatrix defined by the row sets identified in $\mathbf{R}[\cdot, k]$ and the column sets identified in $\mathbf{C}[\cdot, k]$. This general framework subsumes several other integrative decompositions of structured variation. For example, in the unidimensionally linked case with shared columns $J = 1$, JIVE (Lock et al., 2013) identifies components that are either shared across all matrices ($\mathbf{R}[i, k] = 1$ for all i) or specific to a single matrix ($\mathbf{R}[i, k] = 1$ for just one i), while SLIDE (Gaynanova and Li, 2019) allows for components that are shared across any number of matrices. The BIDIFAC method (Park and Lock, 2020) is a special case in which modules are either globally shared ($\mathbf{R}[i, k] = \mathbf{C}[j, k] = 1$ for all i, j), shared across all rows for a given column set ($\mathbf{R}[i, k] = 1$ for all i and $\mathbf{C}[j, k] = 1$ for one j), shared across all columns for a given row set ($\mathbf{C}[j, k] = 1$ for all j and $\mathbf{R}[i, k] = 1$ for one i), or specific to one matrix ($\mathbf{R}[i, k] = 1$ and $\mathbf{C}[j, k] = 1$ for just one i, j).

The BIDIFAC and BIDIFAC+ approaches solve a structured nuclear norm objective:

$$\operatorname{argmin}_{\{\mathbf{S}_{..}^{(k)}\}_{k=1}^K} \frac{1}{2} \|\mathbf{X}_{..} - \sum_{k=1}^K \mathbf{S}_{..}^{(k)}\|_F^2 + \sum_{k=1}^K \lambda_k \|\mathbf{S}_{..}^{(k)}\|_*, \quad (6)$$

subject to $\mathbf{S}_{ij}^{(k)} = \mathbf{0}_{M_i \times N_j}$ if $\mathbf{R}[i, k] = 0$ or $\mathbf{C}[j, k] = 0$. For fixed \mathbf{R} and \mathbf{C} , objective (6) may be solved with an iterative soft-singular value thresholding algorithm that cycles through the estimation of each $\mathbf{S}_{..}^{(k)}$ with applications of Proposition 3. For smaller I and J , \mathbf{R} and \mathbf{C} can enumerate all possible submatrices of the row and column sets; the solution may still be sparse, with $\mathbf{S}_{..}^{(k)} = \mathbf{0}$ if the submatrix corresponding to module k has no particular shared structure. Alternatively, if the number of row and

column sets I and J are large, [Lock et al. \(2022\)](#) describe an algorithm to adaptively estimate \mathbf{R} and \mathbf{C} to give the submatrices with the most significant low-rank structure during estimation.

The choice of the module-specific penalty parameters λ_k is critical. By default, λ_k are selected by an approach motivated by [Proposition 5](#); under the assumption that each matrix \mathbf{X}_{ij} has error variance 1, λ_k depends on the total number of non-zero rows and columns in its submatrix,

$$\lambda_k = \sqrt{\sum_{i=1}^I \mathbf{R}[i, k] M_i} + \sqrt{\sum_{j=1}^J \mathbf{C}[j, k] N_j},$$

for $k = 1, \dots, K$. This approach will shrink to 0 any singular values in the submatrix that are due to error only. In practice each data matrix \mathbf{X}_{ij} is scaled to have error variance 1 beforehand; if the error variance is unknown, then it can be taken as the total variance in \mathbf{X}_{ij} (a conservative choice) or via other approaches such as the median absolute deviation estimator in [Gavish and Donoho \(2017\)](#).

5 Empirical variational BIDIFAC (EV-BIDIFAC)

While the BIDIFAC+ approach has several advantages, one weakness is the required standardization by the noise variance in each data matrix, which is estimated pre-hoc and not within a cohesive framework. More critically, the use of nuclear norm penalties generally results in over-shrinkage of the signal, and the default theoretically-motivated choice of the λ_k 's is particularly conservative. Thus, we aim to develop a more flexible empirical variational Bayes approach to the decomposition of bidimensionally linked matrices, analogous to that for a single matrix in [Section 3](#). Our psuedo-Bayesian model is as follows for $i = 1, \dots, I$ and $j = 1, \dots, J$:

$$\mathbf{X}_{ij} = \sum_{k=1}^K \mathbf{U}_i^{(k)} \mathbf{V}_j^{(k)T} + \mathbf{E}_{i,j} \text{ where } \mathbf{E}_{i,j}[l, m] \stackrel{iid}{\sim} \text{Normal}(0, \sigma_{ij}^2), \quad (7)$$

$$\mathbf{U}_i^{(k)}[\cdot, r] = 0 \text{ if } \mathbf{R}[i, k] = 0 \text{ and } \mathbf{U}_i^{(k)}[\cdot, r] \sim \text{Normal}(\mathbf{0}, \sigma_{ij} \tau_u^2[k, r]) \text{ if } \mathbf{R}[i, k] = 1,$$

$$\mathbf{V}_j^{(k)}[\cdot, r] = 0 \text{ if } \mathbf{C}[j, k] = 0 \text{ and } \mathbf{V}_j^{(k)}[\cdot, r] \sim \text{Normal}(\mathbf{0}, \sigma_{ij} \tau_v^2[k, r]) \text{ if } \mathbf{C}[j, k] = 1.$$

Using repeated applications of [Corollary 4](#), the BIDIFAC+ objective can be shown to give the posterior mode of the model specified in [\(7\)](#) for which $\boldsymbol{\sigma} = \{\sigma_{ij} : i = 1, \dots, I, j = 1, \dots, J\}$ and $\boldsymbol{\tau} = \{\tau_u, \tau_v\}$ are fixed, and $\tau_u^2[k, r] = \tau_v^2[k, r] = 1/\lambda_k$ for all r . We extend this framework by allowing $\boldsymbol{\sigma}$ and $\boldsymbol{\tau}$ to be estimated within a cohesive empirical variational Bayes framework. We approximate the true posterior with a distribution $q(\cdot)$, for which the factor matrices are mutually independent:

$$q(\mathbf{U}, \mathbf{V}) = \prod_{k=1}^K q_{u,k}(\mathbf{U}_k) q_{v,k}(\mathbf{V}_k).$$

The corresponding free energy is given by

$$F(q | \boldsymbol{\sigma}, \boldsymbol{\tau}) = E_q \log \frac{\prod_{k=1}^K q_{u,k}(\mathbf{U}_k) q_{v,k}(\mathbf{V}_k)}{p(\mathbf{X}_{..} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma}) p(\{\mathbf{U}\}_{k=1}^K | \boldsymbol{\tau}_U) p(\{\mathbf{V}\}_{k=1}^K | \boldsymbol{\tau}_V)}, \quad (8)$$

with the densities $p(\cdot)$ in the denominator as specified in model (7). The resulting low-rank fit and decomposition is then given by

$$\hat{\mathbf{S}}_{..} = E_q \left(\sum_{k=1}^K \mathbf{U}_i^{(k)} \mathbf{V}_j^{(k)T} \right) = \sum_{k=1}^K \hat{\mathbf{S}}_{..}^{(k)} \quad (9)$$

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

$$\hat{\mathbf{S}}_{..}^{(k)} = E_{q_{u,k}, q_{v,k}} \left(\mathbf{U}_i^{(k)} \mathbf{V}_j^{(k)T} \right).$$

Algorithm 1 describes an iterative approach to compute the empirical variational Bayes solution. This algorithm converges to a blockwise minimum of the free energy (8) over q and $\boldsymbol{\tau}$ with $\boldsymbol{\sigma}$ fixed at the distinct minimizer of (8) for each \mathbf{X}_{ij} separately. Alternatively, one can update the variances $\boldsymbol{\sigma}$ over the iterative algorithm to minimize free energy under the joint model. In Algorithm 1, the row- and column-set inclusions for the modules, \mathbf{R} and \mathbf{C} , are assumed to be fixed. In practice, for a smaller number linked matrices, one can set \mathbf{R} and \mathbf{C} to enumerate all possible submatrices of shared, partially shared, or specific structure. The estimated decomposition may still be sparse in that $\hat{\mathbf{S}}_{..}^{(k)} = \mathbf{0}$ if module k has no distinct low-rank structure, as demonstrated in Section 8.3. If enumerating all submatrices of the row- and column-sets is not feasible due to a large number matrices, one can specify \mathbf{R} and \mathbf{C} based on prior knowledge. Alternatively, the upper bound on the number distinct modules, K , may be determined beforehand and \mathbf{R} and \mathbf{C} estimated empirically as in Lock et al. (2022).

Algorithm 1 Estimation of EV-BIDIFAC model for fully observed data

1. Estimate σ_{ij} for each \mathbf{X}_{ij} , as in Theorem 7.
 2. Define $\tilde{\mathbf{X}}_{..}$ as in (1), with $\tilde{\mathbf{X}}_{ij} = \mathbf{X}_{ij} / \sigma_{ij}$ for $i = 1, \dots, I$ and $j = 1, \dots, J$.
 3. Initialize $\tilde{\mathbf{S}}_{..}^{(k)} = \mathbf{0}$ for $k = 1, \dots, K$
 4. For $k = 1, \dots, K$:
 - (a) Let $\tilde{\mathbf{X}}_{..}^{(k)}$ be the submatrix of $\tilde{\mathbf{X}}_{..} - \sum_{k' \neq k} \tilde{\mathbf{S}}_{..}^{(k')}$ over row sets indicated by $\mathbf{R}[\cdot, k]$ and column sets $\mathbf{C}[\cdot, k]$.
 - (b) Update the non-zero submatrix in $\tilde{\mathbf{X}}_{..}^{(k)}$ by the application of the singular value thresholding in Theorem 6 to $\tilde{\mathbf{X}}_{..}^{(k)}$.
 5. Repeat Step 3. until convergence of the $\{\tilde{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$.
 6. Unscale the estimated structure: $\hat{\mathbf{S}}_{ij}^{(k)} = \sigma_{ij} \tilde{\mathbf{S}}_{ij}^{(k)}$ for $k = 1, \dots, K$, $j = 1, \dots, J$, $i = 1, \dots, I$.
-

6 Uniqueness

The uniqueness of decompositions into shared, partially shared, and specific low-rank structure requires careful consideration, as the terms in the model are generally not uniquely defined in general without further constraints. For unidimensional decompositions of low-rank structure, conditions of orthogonality are commonly used to assure the different components are uniquely identified (Lock et al., 2013; Feng et al., 2018; Gaynanova and Li, 2019; Yi et al., 2023); however, such conditions are not straightforward to extend to decompositions for bidimensionally linked data. Alternatively, Lock et al. (2022) showed that the terms in the BIDIFAC+ solution are uniquely defined under certain conditions, without orthogonality, as the minimizer of the structured nuclear norm objective in (6). Here we extend this result to a much broader class of penalties, that includes the EV-BIDIFAC decomposition.

For fixed row-set and column-set indicators for the modules, \mathbf{R} and \mathbf{C} , let $\mathbb{S}_{\hat{\mathbf{X}}}$ be the set of possible decompositions that yield a given approximation $\hat{\mathbf{X}}_{..}$:

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

In most setting the cardinality of $|\mathbb{S}_{\hat{\mathbf{X}}}|$ is infinite without further conditions. Define $P_{\boldsymbol{\lambda}}(\cdot)$ for a vector $\boldsymbol{\lambda}$ as a weighted sum of the singular values of a matrix:

$$P_{\boldsymbol{\lambda}}(\mathbf{S}) = \sum \boldsymbol{\lambda}[r] \mathbf{D}_{\mathbf{S}}[r, r],$$

and let $f_{\text{pen}}(\cdot)$ give a structured penalty on the different modules of a decomposition,:

$$f_{\text{pen}}(\{\mathbf{S}_{..}^{(k)}\}_{k=1}^K) = \sum_{k=1}^K P_{\boldsymbol{\lambda}_k}(\mathbf{S}_{..}^{(k)}),$$

with non-zero entries for each $\boldsymbol{\lambda}_k$. Theorem 8 gives sufficient conditions for the minimizer of this penalty to be uniquely determined.

Theorem 8. Consider $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and let $\mathbf{U}_{.}^{(k)} \hat{\mathbf{D}}^{(k)} \mathbf{V}_{.}^{(k)T}$ give the SVD of $\hat{\mathbf{S}}_{..}^{(k)}$ for $k = 1, \dots, K$. The following three properties uniquely identify $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$.

1. $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$ 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 proof of Theorem 8 is given in Appendix A.1. Note that the linear independence conditions are straightforward to verify, and are likely to hold in practice if the ranks of the terms in $\{\hat{\mathbf{S}}_{..}^{(k)}\}_{k=1}^K$ are small relative to the matrix dimensions. As noted in Proposition 9, the terms in any accumulation point of Algorithm 1 also give

a coordinate-wise optimum of $f_{\text{pen}}(\cdot)$ for a certain set of penalties $\{\lambda_k\}_{k=1}^K$, and thus the EV-BIDIFAC solution is unique in this sense.

Proposition 9. *If $\{\hat{\mathbf{S}}_{\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ is an accumulation point of Algorithm 1, then it is a coordinate-wise optimum of $f_{\text{pen}}(\cdot)$ over $\mathbb{S}_{\hat{\mathbf{X}}}$ for some $\{\lambda_k\}_{k=1}^K$.*

7 Missing data imputation

Now we consider the context in which the data are not fully observed and missing data must be imputed. Let \mathcal{M} index observations in the full dataset \mathbf{X}_{\cdot} that are not observed: $\mathcal{M} = \{(m, n) : \mathbf{X}_{\cdot}[m, n] \text{ is missing}\}$. We develop an Expectation-Maximization (EM) approach for imputation that is analogous to previous approaches described for hard-thresholding (Kurucz et al., 2007), soft-thresholding (Mazumder et al., 2010), and BIDIFAC+ (Lock et al., 2022). The procedure for fixed estimates of the noise variance for each matrix, σ_{ij} , is described in Algorithm 2.

Algorithm 2 Estimation of EV-BIDIFAC model for missing data

1. Initialize $\tilde{\mathbf{S}}_{\cdot}^{(k)} = \mathbf{0}$ for $k = 1, \dots, K$
 2. Define $\tilde{\mathbf{X}}_{\cdot}$ as in (1), with $\tilde{\mathbf{X}}_{ij} = \mathbf{X}_{ij}/\sigma_{ij}$ for $i = 1, \dots, I$ and $j = 1, \dots, J$.
 3. Set $\tilde{\mathbf{X}}_{\cdot}[m, n] = \tilde{\mathbf{S}}_{\cdot}[m, n]$ for $(m, n) \in \mathcal{M}$, where $\tilde{\mathbf{S}}_{\cdot} = \sum_{k=1}^K \tilde{\mathbf{S}}_{\cdot}^{(k)}$.
 4. Update $\{\tilde{\mathbf{S}}_{\cdot}^{(k)}\}_{k=1}^K$ as in Step 4. of Algorithm 1.
 5. Repeat Steps 3. and 4. until convergence of the $\{\tilde{\mathbf{S}}_{\cdot}^{(k)}\}_{k=1}^K$.
 6. Unscale the estimated structure: $\hat{\mathbf{S}}_{ij}^{(k)} = \sigma_{ij} \tilde{\mathbf{S}}_{ij}^{(k)}$ for $k = 1, \dots, K$, $j = 1, \dots, J$, $i = 1, \dots, I$.
 7. Impute $\mathbf{X}_{\cdot}[m, n] = \hat{\mathbf{S}}_{\cdot}[m, n]$ for $(m, n) \in \mathcal{M}$.
-

Note that Step 4. in Algorithm 2 is a partial maximization step, as the free energy in (8) is conditionally minimized for each module given the current imputed values, while Step 5. is an expectation step replacing the missing data with their expected value (9) under the current model fit. Further, the algorithm can be shown to minimize the free energy over the missing data $\{\mathbf{X}_{\cdot}[m, n] : (m, n) \in \mathcal{M}\}$; this result is given as Theorem 10, and is appealing because all unknowns (model hyperparameters, parameters and missing entries) are then estimated via optimizing the same unified free energy objective.

Theorem 10. *Algorithm 2 iteratively minimizes the free energy (9) over $q(\cdot)$, τ , and $\{\mathbf{X}_{\cdot}[m, n] : (m, n) \in \mathcal{M}\}$.*

A proof of Theorem 10 is provided in Appendix A.2. For the choice of σ , we suggest two different approaches, depending on the type of missingness. If entries are missing from an entire row or columns of a given data matrix \mathbf{X}_{ij} (i.e., row- or column-wise missing), then σ_{ij} can be estimated by the applying Theorem 2 to the complete submatrix $\{\mathbf{X}_{\cdot}[m, n] : (m, n) \notin \mathcal{M}\}$. If the matrix is missing entries without entire

rows or columns, then we update σ_{ij} iteratively during Algorithm 2 as follows at the end of each cycle: 1. apply Theorem 2 to the matrix \mathbf{X}_{ij} with currently imputed values to obtain σ'_{ij} , 2. update $\sigma_{ij} = \sigma'_{ij} \left(\frac{M_i N_j}{M_i N_j - |\mathcal{M}_{ij}|} \right)$ where $|\mathcal{M}_{ij}|$ is the number of missing entries in \mathbf{X}_{ij} .

8 Simulations

Here we present a collection of simulations to illustrate and compare the proposed approach with respect to various criteria. We explore several conditions covering different signal-to-noise ratios, missing data scenarios, and linked matrix patterns with shared and specific structure. Further simulation results (e.g., exploring different matrix dimensions) are available at <https://ericfrazierlock.com/ev-bidifac-sims.html>.

8.1 Single matrix

We first describe a simple simulation on a single matrix to illustrate the empirical variational Bayes approach to matrix factorization as a flexible compromise between a hard-thresholding and a conservative soft-thresholding approach. We consider a fully observed matrix $\mathbf{X} : 1000 \times 100$, with $\mathbf{X} = c\mathbf{U}\mathbf{V}^T + \mathbf{E}$ where the entries of \mathbf{E} , $\mathbf{U} : M \times 10$ and $\mathbf{V} : N \times 10$ are each independent standard normal, and c is a manipulated value controlling the signal-to-noise ratio. We consider 10 values of c distributed on a log-uniform scale between $c = 0.05$ and $c = 1$, and generate 10 replications of the simulation for each value of c . For each generated dataset, we apply the following approaches:

- OPT: The oracle operator on the singular values when the structure $\mathbf{S} = c\mathbf{U}\mathbf{V}^T$ is known: $\hat{\mathbf{S}}_{\text{OPT}} = \mathbf{U}_X \hat{\mathbf{D}} \mathbf{V}_X^T$ where $\hat{\mathbf{D}}$ is diagonal with entries that solve the OLS problem to minimize $\|\mathbf{S} - \mathbf{U}_X \hat{\mathbf{D}} \mathbf{V}_X^T\|_F^2$.
- EVB: the empirical variational Bayes approach applying Theorems 6 and 7.
- HT: The hard-thresholding approach described in Proposition 1, using the true rank $R = 10$.
- NN: The nuclear norm penalized soft-thresholding approach described in Proposition 2, with $\lambda = \sqrt{M} + \sqrt{N}$ (this correspond to a threshold motivated by Proposition 5 with true noise variance $\sigma = 1$).


For each setting and for each method we compute the relative squared error (RSE) and oracle-normalized squared error (ONSE)

$$\text{RSE} = \frac{\|\mathbf{S} - \hat{\mathbf{S}}\|_F^2}{\|\mathbf{S}\|_F^2} \text{ and } \text{ONSE} = \frac{\|\mathbf{S} - \hat{\mathbf{S}}\|_F^2}{\|\mathbf{S} - \hat{\mathbf{S}}_{\text{OPT}}\|_F^2}. \quad (10)$$

Figure 1 shows the mean values for both RSE and ONSE under the different methods as a function of the signal-to-noise c . The hard-thresholding approach tends to overfit the signal, which can lead to particularly poor recovery when the signal is smaller relative to the noise. In contrast, the nuclear norm penalty is appropriately conservative

in that it never overfits and performs better for lower signal; however, it tends to over-shrink for higher signal and can have several fold higher error than hard-thresholding. The EVB approach is a flexible compromise between the two methods, performing comparably or better than both methods for all scenarios and substantially outperforming them for moderate levels of signal. Moreover, performance of the EVB method is impressively close to the optimal in most scenarios.

Fig. 1 Error in estimating the underlying low-rank signal for a single matrix under different methods, and under different signal-to-noise (s2n) ratios. The left-panel gives relative squared error, and the right-panel gives oracle normalized standard error. All axes are on a log-scale.



We consider another scenario with low-rank structure in which the underlying components have heterogeneous levels of signal. First, we simulate $\mathbf{S}' = \mathbf{U}'\mathbf{V}'^T$ where the entries of $\mathbf{U}' : M \times 10$ and $\mathbf{V}' : 10 \times N$ are independent standard normal. Then, we generate the signal as

$$\mathbf{S} = \mathbf{U}_{S'}\mathbf{D}_S\mathbf{V}_{S'}^T \quad (11)$$


where \mathbf{D}_S is diagonal with entries sampled randomly from a log-uniform distribution between $0.05\sqrt{MN}$ and \sqrt{MN} , and $\mathbf{X} = \mathbf{S} + \mathbf{E}$ with entries of \mathbf{E} independent standard normal. Thus, some singular values of the true low-rank structure \mathbf{S} may correspond to high signal, others may be of moderate signal, and others may be of weak or undetectable signal. We generate 100 replications in this case, and for each we apply the EVB, HT, and NN methods described above, as well as the following additional approaches:

- HT-OPT: the oracle hard-thresholding approach, in which the rank is selected that results in minimal squared error between the estimated signal and the true signal.
- NN-OPT: the oracle soft-thresholding approach, in which the nuclear norm penalty is selected that results in minimal squared error between the estimated signal and the true signal.

- RMT: the low-rank reconstruction approach proposed in [Shabalin and Nobel \(2013\)](#), based on minimizing squared error loss using asymptotic random matrix theory, with the true error variance $\sigma^2 = 1$.

Boxplots showing distribution of RSE and ONSE for recovering the underlying structure across replications are shown in Figure 2. The EVB and RMT methods have performance close to the oracle, and vastly outperform hard-thresholding and nuclear norm penalized approaches. The relatively poor performance of HT-OPT suggests that singular values benefit from some shrinkage, while the poor performance of NN-OPT indicates that the level of shrinkage should not be uniform across singular values.

Fig. 2 Error in estimating underlying low-rank structure in which the rank-1 components have heterogenous signal sizes. The left-panel gives relative squared error (RSE), and the right-panel gives oracle normalized standard error (ONSE). All axes are on a log-scale.

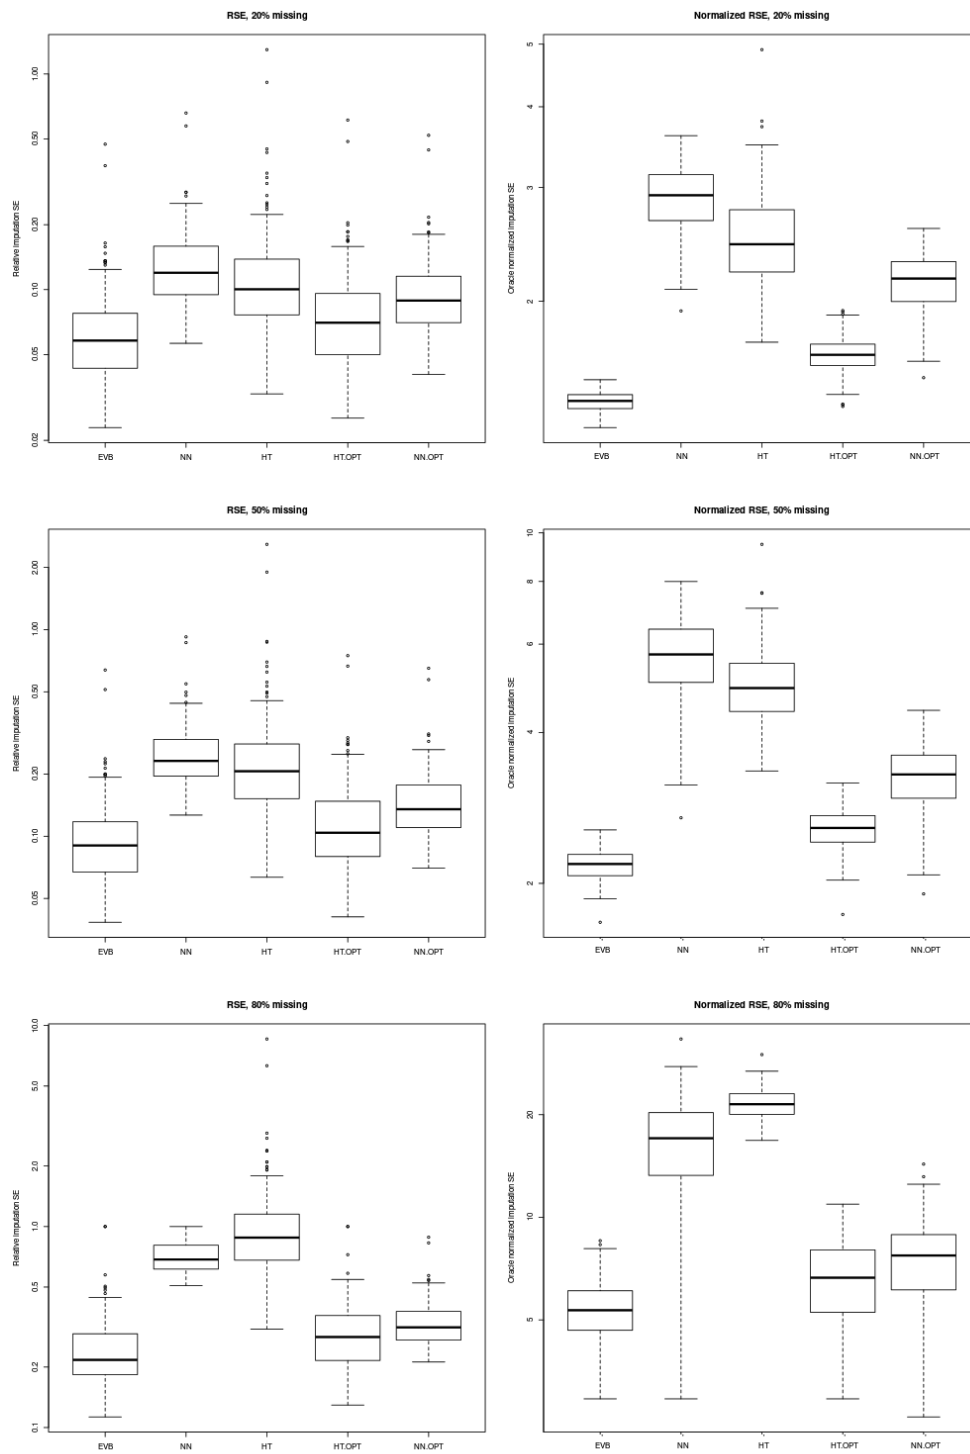


We also compare missing data imputation performance of the EVB approach using Algorithm 2. For each of the 100 simulated datasets generated, we consider scenarios in which 20%, 50%, or 80% of the entries are randomly held out as missing and imputed. In addition to EVB imputation, we apply similar EM-imputation methods for the soft-thresholding (corresponding to *softImpute* of [Mazumder et al. \(2010\)](#)) and hard thresholding ([Kurucz et al., 2007](#)) approaches. The RMT approach is not considered here, as it does not have an associated missing data imputation procedure. For the NN-OPT and HT-OPT approaches, we select the value of λ or R , respectively, that minimizes squared imputation error for the signal of the missing data

$$\text{RSE}_{\text{miss}} = \frac{\sum_{(m,n) \in \mathcal{M}} (\mathbf{S}[m,n] - \hat{\mathbf{S}}[m,n])^2}{\sum_{(m,n) \in \mathcal{M}} \mathbf{S}[m,n]^2}. \quad (12)$$

For oracle normalized imputation error $\text{ONSE}_{\text{miss}}$, we divide RSE_{miss} by the RSE for the OPT method in (10). The results are shown in Figure 3, and demonstrate that the

Fig. 3 Missing data imputation accuracy for different levels of missingness. The left column gives RSE_{miss} and the right gives $ONSE_{miss}$.



EVB imputation approach is stable for different levels of missingness and substantially outperforms the other approaches.

8.2 Two linked matrices

Here we generate two matrices, $\mathbf{X}_1 : 500 \times 100$ and $\mathbf{X}_2 : 500 \times 100$, that are linked by columns: $\mathbf{X} = [\mathbf{X}_1; \mathbf{X}_2] : 1000 \times 100$. We consider this context primarily to allow comparison of the EV-BIDIFAC approach with several other methods that identify shared and unshared low-rank structure in this setting.

First, we perform an illustrative simulation akin to that in Figure 1 for a single matrix. We simulate shared structure $\mathbf{S}_0 = c\mathbf{U}_0\mathbf{V}_0^T$ and individual structures $\mathbf{S}_1 = c\mathbf{U}_1\mathbf{V}_1^T$, $\mathbf{S}_2 = c\mathbf{U}_2\mathbf{V}_2^T$, where $\mathbf{U}_0 : 1000 \times 5$, $\mathbf{V}_0 : 100 \times 5$, $\mathbf{U}_1 : 500 \times 5$, $\mathbf{U}_2 : 500 \times 5$, $\mathbf{V}_1 : 100 \times 5$, and $\mathbf{V}_2 : 100 \times 5$ each have independent entries from a standard normal distribution. The observed data are then generated as $\mathbf{X} = \mathbf{S}_0 + [\mathbf{S}_1; \mathbf{S}_2] + \mathbf{E}$ where entries of \mathbf{E} are independent standard normal. Thus c again controls the signal to noise ratio and we consider 10 values of c distributed on a log-uniform scale between $c = 0.05$ and $c = 1$, and generate 10 replications for each value of c . Here we consider estimates for the following methods:

- EVB: the EV-BIDIFAC approach described in Algorithm 2
- JIVE: the JIVE method implemented in O’Connell and Lock (2016) for true shared and unshared ranks $R = 5, R_1 = 5, R_2 = 5$.
- UNIFAC: the BIDIFAC method (Park and Lock, 2020) for vertical integration using the true noise variance $\sigma^2 = 1$.


Here JIVE is analogous to HT and UNIFAC to NN for a single matrix, as they involve iterative applications HT or NN, respectively, to the shared and unshared components. For each method we compute RSE for the overall structure $\hat{\mathbf{S}} = \hat{\mathbf{S}}_0 + [\hat{\mathbf{S}}_1; \hat{\mathbf{S}}_2]$ as in (10), and also compute the relative decomposition squared error (RDSE) as

$$\text{RDSE} = \frac{\|\mathbf{S}_0 - \hat{\mathbf{S}}_0\|_F^2 + \|\mathbf{S}_1 - \hat{\mathbf{S}}_1\|_F^2 + \|\mathbf{S}_2 - \hat{\mathbf{S}}_2\|_F^2}{\|\mathbf{S}_0\|_F^2 + \|\mathbf{S}_1\|_F^2 + \|\mathbf{S}_2\|_F^2}.$$

The results are shown in Figure 4. The results for overall RSE are similar to that for a single matrix, with UNIFAC appropriately conservative for lower signal but substantially underperforming JIVE for higher signal due to over-shrinkage. However, per the RDSE results, UNIFAC can still identify the shared and unshared components more accurately with higher signal – this phenomenon was noted previously (Lock et al., 2022) and is in part due to orthogonality restrictions imposed by JIVE and related methods that are required for uniqueness of the decomposition without shrinkage. The EVB approach again acts as a flexible compromise between the two methods, and substantially outperforms both for many s2n scenarios, particularly with respect to RDSE.

Now, we compare RSE and RDSE under a scenario with heterogenous signal levels for the factors in each of the shared and unshared structures. The rank-5 signal matrices $\mathbf{S}_0, \mathbf{S}_1$, and \mathbf{S}_2 are each generated as in (11) with singular values again generated uniformly from a log-uniform distribution between $0.05\sqrt{MN} * 0.05$ and \sqrt{MN} for

Fig. 4 Error in estimating the underlying low-rank signal for two linked matrices under different signal-to-noise (s2n) ratios. The left-panel gives the RSE, and the right-panel gives RDSE. All axes are on a log-scale.



shared structure and between $0.05\sqrt{M_iN}$ and $\sqrt{M_iN}$ for specific structure, and 100 replications are generated for this scenario. In addition to the EVB, JIVE, UNIFAC methods we also consider the following approaches:

- SLIDE: the structural learning and integrative decomposition (SLIDE) method (Gaynanova and Li, 2019), with the true number of shared and unshared components (i.e., the ranks) specified.
- HNN.OPT: the hierarchical nuclear norm (HNN) approach (Yi et al., 2023), with tuning parameters selected on a grid to minimize actual RSE for the true signal.

Figure 5 gives distribution of RSE and RDSE over the replications for the different methods. The EVB approach universally outperforms other methods in terms of both RSE and RDSE, and the benefit is large (by a factor of 2 or more) for both criteria in several cases. While the HNN method is motivated by a model with shared and unshared low-rank structure, it does not provide an explicit decomposition of the structure and thus it is not shown for RDSE.


8.3 Bidimensionally linked matrices

Now, we consider the full bidimensionally linked framework in (1), with $I = J = 2$, $N_1 = N_2 = 50$ and $D_1 = D_2 = 500$. Data were generated as

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

where the entries of \mathbb{E} are independent standard normal. There are $K = 9$ possible distinct modules of low-rank structure in this setting: globally shared, shared across

Fig. 5 RSE (left) and RDSE (right) for low-rank structure with heterogenous signal levels for two linked matrices.



rows for either of the of the $J = 2$ column sets, shared across columns for either of the $I = 2$ row sets, or structure specific to any of the 4 matrices \mathbf{X}_{ij} . We generate 100 datasets, where for each replication 5 of the 9 possible modules are randomly selected to have non-zero low-rank structure and the others are set to zero, $\mathbf{S}_{\cdot\cdot}^{(k)} = \mathbf{0}$. All non-zero modules are generated to have rank 2 structure on the given submatrix as in (11). For each simulated dataset, we apply the following methods:

- EB-BIDI: the approach in Algorithm 1, with \mathbf{R} and \mathbf{C} enumerating all $K = 9$ possible modules.
- BIDIFAC: the BIDIFAC algorithm described in (Lock et al., 2022) with \mathbf{R} and \mathbf{C} enumerating all possible modules and using the true noise variance $\sigma^2 = 1$.
- EB-SEP: the EVB approach applied to each matrix \mathbf{X}_{ij} separately
- EB-JOINT: the EVB approach applied to the overall matrix $\mathbf{X}_{\cdot\cdot}$.


For each method, we compute the overall RSE for $\mathbf{S}_{\cdot\cdot} = \sum_{k=1}^K \mathbf{S}_{\cdot\cdot}^{(k)}$ as in (10), and for EB-BIDI and BIDIFAC we compute the RDSE as

$$\text{RDSE} = \frac{\sum_{k=1}^9 \|\mathbf{S}_{\cdot\cdot}^{(k)} - \hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\|_F^2}{\sum_{k=1}^9 \|\mathbf{S}_{\cdot\cdot}^{(k)}\|_F^2}.$$

The results are summarized in Figure 6. Overall RSE is substantially better for EB-BIDI vs. other methods; in comparison to EB-SEP and EB-JOINT, this illustrates the advantages of a parsimonious decomposition of the low-rank signal into the different shared and unshared components with respect recovering the overall signals. Moreover, EB-BIDI has substantially better RDSE compared to BIDIFAC, demonstrating the advantages of the empirical variational Bayes objective with respect to accurately decomposing the structure into its different components. Moreover, EB-BIDI correctly estimated modules with no true structure ($\mathbf{S}_{\cdot\cdot}^{(k)} = \mathbf{0}$) as $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}$ for all of the 400 total

instances (100 total datasets with 4 possible modules set to $\mathbf{0}$ in each case); on the other hand, it estimated $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{0}$ in only $19/500 = 3.8\%$ of instances when $\mathbf{S}_{\cdot\cdot}^{(k)} \neq \mathbf{0}$, which occurred when $\mathbf{S}_{\cdot\cdot}^{(k)}$ was difficult to detect due to a small signal size. This demonstrates how the approach can correctly recover the underlying rank-sparsity structure, only estimating non-zero structure for a module if it truly has distinct low-rank structure on its given rows and columns.

Fig. 6 RSE (left) and RDSE (right) for the scenario with 2×2 bidimensionally linked matrices.




We extend this simulation design to consider missing data imputation accuracy for the generated datasets. We consider two different missingness scenarios: an *entry-wise missing* scenario in which 20% of the entries are randomly set to missing in each matrix, and a *blockwise missing* scenario in which 10% of the columns and 10% of the rows are randomly selected to have all of their entries missing in each matrix. In addition to imputation using the previous methods, we also consider the NN.OPT and HT.OPT methods as described in Section 8.1. For entrywise imputation the RSE_{miss} is computed as in (12). For blockwise imputation, the true underlying structure is taken to be the sum of column-shared structure for missing rows and the sum of row-shared structure for missing columns, as this is the only estimable signal when an entire row or column is missing. That is, letting $\mathbf{S}_{\cdot\cdot}^C = \sum_{k:\mathbf{C}[\cdot,k]=\mathbf{1}} \mathbf{S}_{\cdot\cdot}^{(k)}$ and $\mathbf{S}_{\cdot\cdot}^R = \sum_{k:\mathbf{R}[\cdot,k]=\mathbf{1}} \mathbf{S}_{\cdot\cdot}^{(k)}$,

$$\text{RSE}_{\text{miss}} = \frac{\sum_{(m,n) \in \mathcal{M}^R} (\mathbf{S}_{\cdot\cdot}^C[m,n] - \hat{\mathbf{S}}_{\cdot\cdot}[m,n])^2 + \sum_{(m,n) \in \mathcal{M}^C} (\mathbf{S}_{\cdot\cdot}^R[m,n] - \hat{\mathbf{S}}_{\cdot\cdot}[m,n])^2}{\sum_{(m,n) \in \mathcal{M}^R} \mathbf{S}_{\cdot\cdot}^C[m,n]^2 + \sum_{(m,n) \in \mathcal{M}^C} \mathbf{S}_{\cdot\cdot}^R[m,n]^2}, \quad (13)$$

where \mathcal{M}^R indexes row-missing entries and \mathcal{M}^C indexes column missing entries. The results are shown in Figure 7. The EB-BIDI method substantially and consistently outperforms the other approaches with respect to missing data imputation, and the

advantages are particularly large for blockwise missing imputation. Note that the EB-SEP approach does not estimate any structure for blockwise missing data, as it is not estimable when considering each matrix separately.

Fig. 7 RSE for entrywise missing data imputation (left) and blockwise missing data imputation (right), for the scenario with 2×2 bidimensionally linked matrices.

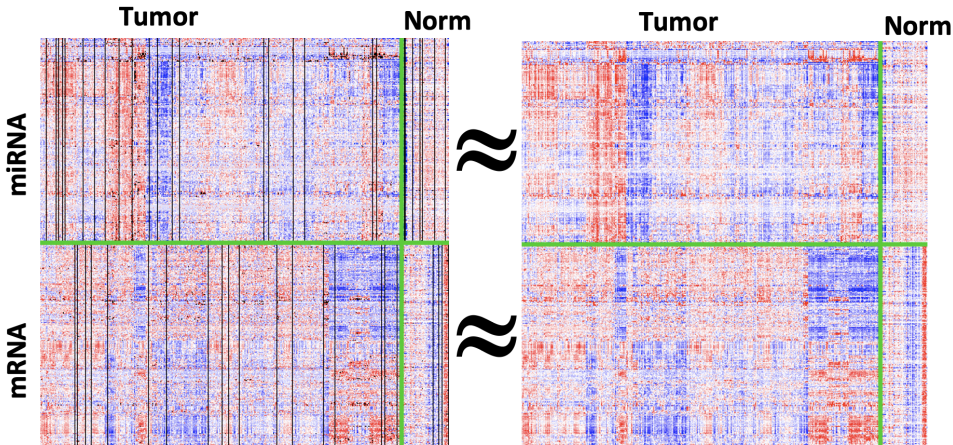


9 Application to BRCA Data

Here we describe an application to gene expression (mRNA) and microRNA (miRNA) data for breast cancer (BRCA) from the Cancer Genome Atlas (TCGA) (TCGA Research Network et al., 2012). The data were processed as described in Park and Lock (2020), resulting in $D_1 = 500$ mRNAs and $D_2 = 500$ miRNAs for $N_1 = 660$ cancer tumor samples and $N_2 = 86$ non-cancerous breast tissue samples from unrelated individuals. Thus, the data can be represented as bidimensionally linked matrices as in (1), where \mathbf{X}_{11} is mRNA data for cancer samples, \mathbf{X}_{12} is mRNA data for normal samples, \mathbf{X}_{21} is miRNA data for cancer samples and \mathbf{X}_{22} is miRNA data for normal samples. The decomposition of data into different modules of shared structured variation is very well-motivated in this context, as low-rank structure on any of the $K = 9$ possible modules is plausible and interpretable. For example, it is reasonable to expect shared structure across mRNA and miRNAs, as miRNAs can regulate mRNA. It is also reasonable to expect some shared structure across tumor and normal samples, as breast tumors arise from normal tissue and so they likely share some patterns of molecular variation. However, it is also reasonable to expect specific structure in each case, e.g., patterns of molecular variability that are ‘cancer-specific’ (i.e., not present in normal tissue) are of particular interest.

We applied EV-BIDIFAC wherein \mathbf{R} and \mathbf{C} enumerate all $K = 9$ possible modules, and significant distinct low-rank structure was identified for each module ($\hat{\mathbf{S}}_{..}^{(k)} \neq \mathbf{0}$

Fig. 8 Heatmaps of the BRCA data (left) and the full low-rank structure estimated by EV-BIDIFAC. Higher values are colored red, lower values are colored blue, and missing columns are colored black.



for all k). Figure 8 shows heatmaps of the original data (including missing data for miRNA and mRNA) and the resulting overall fit $\hat{\mathbf{S}}_{..} = \sum_{k=1}^K \hat{\mathbf{S}}_{..}^{(k)}$. While all estimated structures are low-rank, it captures the original data well and imputation of the missing columns appears reasonable.

We further compare imputation accuracy for data under entry-wise, column-wise or row-wise missingness under a cross-validation scheme. The accuracy of entry-wise imputation is of interest because it is a good proxy for the accuracy of the the overall fit of the estimated structure. Column-wise imputation accuracy is of interest for two reasons: (1) the actual data have column-wise missingness, as not all samples have both mRNA and miRNA data in the full TCGA cohort, and (2) it provides a robust measure of the level of association between the two data sources, addressing how well can we predict gene expression with miRNA data, and vice versa. Row-wise imputation accuracy is also of interest, as it provides a measure of similarity in molecular structure between cancer and normal samples (e.g., how well can we predict levels of a missing gene in cancer samples, given its associations in normal samples).

We create 20 folds in the data, where for each fold 5% of the columns, 5% of the rows, and 5% of the remaining entries in each matrix are randomly withheld as missing. For each fold, the missing data are imputed using the EV-BIDIFAC, BIDIFAC, EB-SEP or EB-JOINT approaches described in Section 8.3. In each case, the relative squared error for the true observed data is computed and averaged across the four datasets yielding

$$\text{MRSE}_{\text{miss}} = \frac{1}{4} \sum_{i=1}^2 \sum_{j=1}^2 \frac{\sum_{(m,n) \in \mathcal{M}_{ij}} (\mathbf{X}_{ij}[m, n] - \hat{\mathbf{S}}_{ij}[m, n])^2}{\sum_{(m,n) \in \mathcal{M}_{ij}} \mathbf{X}_{ij}[m, n]^2}.$$

The $\text{MRSE}_{\text{miss}}$ under different imputation scenarios are summarized in Table 1. The EV-BIDIFAC approach performs relatively well across all scenarios. For entry-wise missing data EV-BIDIFAC performs much better than BIDIFAC (likely due to over-shrinkage of the signal for BIDIFAC) and similarly to EB-SEP, indicating that structure can be well-estimated by considering each dataset separately; however, separate estimation does not allow for row- or column imputation and does not estimate the shared variation that is of interest. The $\text{MRSE}_{\text{miss}}$ for row-missing data is relatively modest (0.90) compared to entry-wise missing (0.38), indicating that there is little shared information between tumor and normal cells. The joint approach EB-JOINT on $\mathbf{X}_{..}$ performs relatively poorly across all scenarios as it tends to overfit shared variation when there substantial specific variation across platforms (mRNA and miRNA) and cohorts (tumor and normal); this illustrates the benefits of a comprehensive decomposition of shared and specific variation across the constituent matrices.

Table 1 Imputation accuracy ($\text{MRSE}_{\text{miss}}$) for BRCA and normal tissue mRNA and miRNA data.

Method Name	Entry-missing	Column-missing	Row-missing	Overall
EV-BIDIFAC	0.389	0.756	0.903	0.682
BIDIFAC	0.526	0.870	0.909	0.768
EB-SEP	0.381	1.00	1.00	0.79
EB-JOINT	0.510	1.01	1.67	1.06

10 Conclusion and discussion

The proposed EV-BIDIFAC approach can be considered as a flexible compromise between hard- and soft- thresholding techniques for the linked matrix decomposition problem, and has impressive performance with respect to 1. estimating underlying low-rank structure, 2. decomposing underlying structure into its shared, partially-shared and specific components, and 3. missing data imputation under a variety of scenarios. An appealing aspect of the method is that it derives from a unified and intuitive model with a corresponding objective function that does not require any pre-hoc choices or parameters to tune. In simulation comparisons the proposed approach tended to outperform other methods even when their parameters are set to the true values (e.g., true rank(s) or true error variance) or ideal values (e.g., to minimize error with knowledge of the true structure).

While we have focused on obtaining a point estimate for underlying structure that minimizes free energy, this could also be used to empirically estimate hyperparameters and initialize a sampling algorithm for the associated fully Bayesian model. This would capture uncertainty in the underlying factorization, which can be used for multiple imputation of missing values or to propagate uncertainty if the lower-dimensional components are used in subsequent statistical modeling. A related Bayesian model

is described in [Samorodnitsky et al. \(2022\)](#), which also illustrates the importance of uncertainty propagation.

Supplementary materials

Reproducible R scripts for all results presented in this manuscript are available in a supplementary zipped folder. Software to perform the EV-BIDIFAC method is available at <https://github.com/lockEF/bidifac>. Additional simulation study results are available at <https://ericfrazierlock.com/ev-bidifac-sims.html>.

Acknowledgement

This work was supported by the NIH National Institute of General Medical Sciences (NIGMS) grant R01-GM130622.

Appendix A Proofs

A.1 Proof of Theorem 8

The proof for Theorem 8 follows a similar structure to that for Theorem 1 in [Lock et al. \(2022\)](#). We first establish four lemmas, Lemmas 11, 12, 13, and 14, which are useful to prove the main result.

Lemma 11. *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 penalty for given $\{\boldsymbol{\lambda}\}_{k=1}^K$:*

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

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 convexity of $P_{\boldsymbol{\lambda}}$, which implies that for any two matrices of equal size $\hat{\mathbf{A}}$ and $\tilde{\mathbf{A}}$,

$$P_{\boldsymbol{\lambda}}(\alpha \hat{\mathbf{A}} + (1 - \alpha) \tilde{\mathbf{A}}) \leq \alpha P_{\boldsymbol{\lambda}}(\hat{\mathbf{A}}) + (1 - \alpha) P_{\boldsymbol{\lambda}}(\tilde{\mathbf{A}}). \quad (\text{A1})$$

Applying (A1) to each additive term in f_{pen} gives

$$\begin{aligned} f_{\text{pen}}\left(\{\alpha\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + (1-\alpha)\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K\right) &\leq \alpha f_{\text{pen}}(\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) + (1-\alpha)f_{\text{pen}}(\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K) \\ &= \min_{\mathbb{S}_{\hat{\mathbf{X}}}} f_{\text{pen}}(\{\mathbf{S}_{\cdot\cdot}^{(k)}\}_{k=1}^K). \end{aligned} \quad (\text{A2})$$

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 (A2) must be an equality, and it follows that the inequality (A1) must be an equality for each penalized term in the decomposition. \square

Lemma 12. *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 $\tilde{\mathbf{A}} = \tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T$ where $\tilde{\mathbf{D}}$ is diagonal and $\|\tilde{\mathbf{A}}\|_* = \|\tilde{\mathbf{D}}\|_*$.*

Proof. This result is proved in Lemma 2 of Lock et al. (2022). \square

Lemma 13. *Take two matrices $\hat{\mathbf{A}}$ and $\tilde{\mathbf{A}}$. If $P_{\lambda}(\hat{\mathbf{A}} + \tilde{\mathbf{A}}) = P_{\lambda}(\hat{\mathbf{A}}) + P_{\lambda}(\tilde{\mathbf{A}})$, 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 $P_{\lambda}(\hat{\mathbf{A}}) = P_{\lambda}(\hat{\mathbf{D}})$, and $\tilde{\mathbf{A}} = \tilde{\mathbf{U}}\tilde{\mathbf{D}}\tilde{\mathbf{V}}^T$ where $\tilde{\mathbf{D}}$ is diagonal and $P_{\lambda}(\tilde{\mathbf{A}}) = P_{\lambda}(\tilde{\mathbf{D}})$.*

Proof. The result follows from repeated applications of Lemma 12 to the rank-1 terms of $\hat{\mathbf{A}}$ and $\tilde{\mathbf{A}}$. Note that $P_{\lambda}(\hat{\mathbf{A}} + \tilde{\mathbf{A}}) = P_{\lambda}(\hat{\mathbf{A}}) + P_{\lambda}(\tilde{\mathbf{A}})$ implies

$$P_{\lambda}\left(\sum_r \mathbf{U}_{\hat{\mathbf{A}}}[\cdot, r]\mathbf{D}_{\hat{\mathbf{A}}}[r, r]\mathbf{V}_{\hat{\mathbf{A}}}[\cdot, r]^T + \mathbf{U}_{\tilde{\mathbf{A}}}[\cdot, r]\mathbf{D}_{\tilde{\mathbf{A}}}[r, r]\mathbf{V}_{\tilde{\mathbf{A}}}[\cdot, r]^T\right) = \sum_r \lambda[r](\mathbf{D}_{\hat{\mathbf{A}}}[r, r] + \mathbf{D}_{\tilde{\mathbf{A}}}[r, r]),$$

which requires that

$$\begin{aligned} &\|\mathbf{U}_{\hat{\mathbf{A}}}[\cdot, r]\mathbf{D}_{\hat{\mathbf{A}}}[r, r]\mathbf{V}_{\hat{\mathbf{A}}}[\cdot, r]^T + \mathbf{U}_{\tilde{\mathbf{A}}}[\cdot, r]\mathbf{D}_{\tilde{\mathbf{A}}}[r, r]\mathbf{V}_{\tilde{\mathbf{A}}}[\cdot, r]^T\|_* \\ &= \|\mathbf{U}_{\hat{\mathbf{A}}}[\cdot, r]\mathbf{D}_{\hat{\mathbf{A}}}[r, r]\mathbf{V}_{\hat{\mathbf{A}}}[\cdot, r]^T\|_* + \|\mathbf{U}_{\tilde{\mathbf{A}}}[\cdot, r]\mathbf{D}_{\tilde{\mathbf{A}}}[r, r]\mathbf{V}_{\tilde{\mathbf{A}}}[\cdot, r]^T\|_* \end{aligned}$$

for each r . \square

Lemma 14. *Take two decompositions $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$ and $\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K \in \mathbb{S}_{\hat{\mathbf{X}}}$, and assume that both satisfy*

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

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

Proof. The result follows as a direct corollary of Lemmas 11 and 13, as Lemma 11 implies $P_{\lambda}(\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} + \tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}) = P_{\lambda}(\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}) + P_{\lambda}(\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)})$ for each k , and then Lemma 13 implies the result. \square

Theorem 8 is then established as follows:

Proof. Take two decomposition $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$ and $\{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$ that satisfy properties 1., 2., and 3. of Theorem 1; we will show that $\{\hat{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K = \{\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)}\}_{k=1}^K$. For each $k = 1, \dots, K$, write $\hat{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{U}_{\cdot\cdot}^{(k)} \hat{\mathbf{D}}^{(k)} \mathbf{V}_{\cdot\cdot}^{(k)T}$ and $\tilde{\mathbf{S}}_{\cdot\cdot}^{(k)} = \mathbf{U}_{\cdot\cdot}^{(k)} \tilde{\mathbf{D}}^{(k)} \mathbf{V}_{\cdot\cdot}^{(k)T}$ as in Lemma 14. 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 $\mathbf{R}[\cdot, 1] = [1 \ 0 \ \dots \ 0]^T$ and $\mathbf{C}[\cdot, 1] = [1 \ 0 \ \dots \ 0]^T$. By way of contradiction, assume $\hat{\mathbf{D}}^{(1)}[1, 1] > 0$ and $\tilde{\mathbf{D}}^{(1)}[1, 1] = 0$. The linear independence of $\{\mathbf{V}_j^{(k)}[\cdot, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0\}$ and $\{\mathbf{V}_j^{(k)}[\cdot, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$ implies that

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

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

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

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

$$\mathbf{U}^{(1)}[\cdot, 1] \in \text{span}\{\mathbf{U}_{\cdot\cdot}^{(k)}[\cdot, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0, k > 1, \text{ and } \mathbf{R}[i, k] = 0 \text{ for any } i > 1\}. \quad (\text{A3})$$

Note that (A3) implies $\mathbf{U}_1^{(1)}[\cdot, 1] \in \text{row}(\mathbf{X}_{12} + \dots + \text{row}(\mathbf{X}_{1J}))$, however, this is contradicted by the linear independence of $\mathbf{U}_1^{(1)}[\cdot, 1]$ and $\{\mathbf{U}_i^{(k)}[\cdot, r] : \tilde{\mathbf{D}}^{(k)}[r, r] > 0, k > 1\}$. Thus, we conclude that $\tilde{\mathbf{D}}^{(1)}[1, 1] > 0$ implies $\hat{\mathbf{D}}^{(1)}[1, 1] > 0$. 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 $\{\mathbf{U}_i^{(k)}[\cdot, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0 \text{ or } \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $i = 1, \dots, I$, and $\{\mathbf{V}_j^{(k)}[\cdot, r] : \hat{\mathbf{D}}^{(k)}[r, r] > 0 \text{ or } \tilde{\mathbf{D}}^{(k)}[r, r] > 0\}$ are linearly independent for $j = 1, \dots, J$. Thus,

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

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

A.2 Proof of Theorem 10

It is clear that the partial ‘M-steps’ in Algorithm 2 each partially optimize the free energy. It remains to be shown that the missing entries minimize free energy when they are replaced with their conditional expectation under the model. Note that the free energy (8) can be expressed as

$$E_q \log \left(\prod_{k=1}^K q_{u,k}(\mathbf{U}_k) q_{v,k}(\mathbf{V}_k) \right) - E_q \log (p(\{\mathbf{U}\}_{k=1}^K | \boldsymbol{\tau}_{\mathbf{U}}) p(\{\mathbf{V}\}_{k=1}^K | \boldsymbol{\tau}_{\mathbf{V}})) - E_q \log (p(\{\mathbf{X}_{..}\} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma})).$$

Consider the last term in the expression, decomposing the missing and non-missing entries of $\mathbf{X}_{..}$:

$$\begin{aligned} E_q \log (p(\mathbf{X}_{..} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma})) &= E_q \log (p(\{\mathbf{X}_{..}[m, n] : (m, n) \notin \mathcal{M}\} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma}) p(\{\mathbf{X}_{..}[m, n] : (m, n) \in \mathcal{M}\} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma})) \\ &= E_q \log (p(\{\mathbf{X}_{..}[m, n] : (m, n) \notin \mathcal{M}\} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma})) + \quad (\text{A4}) \\ &\quad E_q \log (p(\{\mathbf{X}_{..}[m, n] : (m, n) \in \mathcal{M}\} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma})). \end{aligned}$$

Missing data $\{\mathbf{X}_{..}[m, n] : (m, n) \in \mathcal{M}\}$ are random variables that are independent and normally distributed given $\{\mathbf{U}, \mathbf{V}\}$ as in (7). Thus, the last term in (A4) can be expressed as follows:

$$\begin{aligned} E_q \log (p(\{\mathbf{X}_{..}[m, n] : (m, n) \in \mathcal{M}\} | \mathbf{U}, \mathbf{V}, \boldsymbol{\sigma})) &= c_1 + \sum_{(m, n) \in \mathcal{M}} \frac{1}{2\boldsymbol{\sigma}^2[m, n]} E_q (\mathbf{X}_{..}[m, n] - \mathbf{S}_{..}[m, n])^2 \\ &\geq c_1 + \sum_{(m, n) \in \mathcal{M}} \frac{1}{2\boldsymbol{\sigma}^2[m, n]} E_q (\mathbf{S}_{..}[m, n] - \mathbf{S}_{..}[m, n])^2. \end{aligned}$$

Thus, for any given q , the free energy is minimized at $\mathbf{X}_{..}[m, n] = E_q \mathbf{S}_{..}[m, n] = \hat{\mathbf{S}}_{..}[m, n]$ for $(m, n) \in \mathcal{M}$.

References

- Attias, H.: Inferring parameters and structure of latent variable models by variational Bayes. In: Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, pp. 21–30. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA (1999)
- Feng, Q., Jiang, M., Hannig, J., Marron, J.: Angle-based joint and individual variation explained. *Journal of multivariate analysis* **166**, 241–265 (2018)
- Fox, C.W., Roberts, S.J.: A tutorial on variational Bayesian inference. *Artificial intelligence review* **38**, 85–95 (2012)
- Gavish, M., Donoho, D.L.: Optimal shrinkage of singular values. *IEEE Transactions on Information Theory* **63**(4), 2137–2152 (2017)

- Gaynanova, I., Li, G.: Structural learning and integrative decomposition of multi-view data. *Biometrics* **75**(4), 1121–1132 (2019)
- Josse, J., Sardy, S.: Adaptive shrinkage of singular values. *Statistics and Computing* **26**, 715–724 (2016)
- Kurucz, M., Benczúr, A.A., Csalogány, K.: Methods for large scale SVD with missing values. In: *Proceedings of KDD Cup and Workshop*, vol. 12, pp. 31–38 (2007)
- Lock, E.F., Hoadley, K., Marron, J.S., Nobel, A.B.: Joint and Individual Variation Explained (JIVE) for integrated analysis of multiple data types. *The Annals of Applied Statistics* **7**(1), 523 (2013)
- Lock, E.F., Park, J.Y., Hoadley, K.: Bidimensional linked matrix factorization for pan-omics pan-cancer analysis. *Annals of Applied Statistics* **16**(1), 193–215 (2022)
- Mazumder, R., Hastie, T., Tibshirani, R.: Spectral regularization algorithms for learning large incomplete matrices. *The Journal of Machine Learning Research* **11**, 2287–2322 (2010)
- Nakajima, S., Sugiyama, M.: Analysis of empirical map and empirical partially bayes: Can they be alternatives to variational Bayes? In: *Artificial Intelligence and Statistics*, pp. 20–28 (2014). PMLR
- Nakajima, S., Sugiyama, M., Babacan, S.D., Tomioka, R.: Global analytic solution of fully-observed variational Bayesian matrix factorization. *Journal of Machine Learning Research* **14**(1), 1–37 (2013)
- Nakajima, S., Tomioka, R., Sugiyama, M., Babacan, S.: Perfect dimensionality recovery by variational Bayesian pca. In: Pereira, F., Burges, C.J., Bottou, L., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*, vol. 25 (2012)
- O’Connell, M.J., Lock, E.F.: R.JIVE for exploration of multi-source molecular data. *Bioinformatics* **32**(18), 2877–2879 (2016)
- Park, J.Y., Lock, E.F.: Integrative factorization of bidimensionally linked matrices. *Biometrics* **76**(1), 61–74 (2020)
- Rudelson, M., Vershynin, R.: 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*, pp. 1576–1602 (2010). World Scientific
- Shabalin, A.A., Nobel, A.B.: Reconstruction of a low-rank matrix in the presence of Gaussian noise. *Journal of Multivariate Analysis* **118**, 67–76 (2013)

- Schouteden, M., Van Deun, K., Wilderjans, T.F., Van Mechelen, I.: Performing discosca to search for distinctive and common information in linked data. *Behavior research methods* **46**(2), 576–587 (2014)
- Samorodnitsky, S., Wendt, C.H., Lock, E.F.: Bayesian simultaneous factorization and prediction using multi-omic data. *arXiv preprint arXiv:2211.16403* (2022)
- TCGA Research Network, *et al.*: Comprehensive molecular portraits of human breast tumors. *Nature* **490**(7418), 61 (2012)
- Tang, T.M., Allen, G.I.: Integrated principal components analysis. *Journal of Machine Learning Research* **22**(198), 1–71 (2021)
- Wang, W., Stephens, M.: Empirical Bayes matrix factorization. *Journal of Machine Learning Research* **22**(120), 1–40 (2021)
- Yang, Z., Michailidis, G.: A non-negative matrix factorization method for detecting modules in heterogeneous omics multi-modal data. *Bioinformatics* **32**(1), 1–8 (2016)
- Yi, S., Wong, R.K.W., Gaynanova, I.: Hierarchical nuclear norm penalization for multi-view data integration. *Biometrics* (2023)