

Empirical Bayes Matrix Factorization

Wei Wang

*Department of Statistics
University of Chicago
Chicago, IL, USA*

WEIWANG@STATISTICS.UCHICAGO.EDU

Matthew Stephens

*Department of Statistics and Department of Human Genetics
University of Chicago
Chicago, IL, USA*

MSTEPHENS@UCHICAGO.EDU

Editor: Sayan Mukherjee

Abstract

Matrix factorization methods, which include Factor analysis (FA) and Principal Components Analysis (PCA), are widely used for inferring and summarizing structure in multivariate data. Many such methods use a penalty or prior distribution to achieve sparse representations (“Sparse FA/PCA”), and a key question is how much sparsity to induce. Here we introduce a general Empirical Bayes approach to matrix factorization (EBMF), whose key feature is that it estimates the appropriate amount of sparsity by estimating prior distributions from the observed data. The approach is very flexible: it allows for a wide range of different prior families and allows that each component of the matrix factorization may exhibit a different amount of sparsity. The key to this flexibility is the use of a variational approximation, which we show effectively reduces fitting the EBMF model to solving a simpler problem, the so-called “normal means” problem. We demonstrate the benefits of EBMF with sparse priors through both numerical comparisons with competing methods and through analysis of data from the GTEx (Genotype Tissue Expression) project on genetic associations across 44 human tissues. In numerical comparisons EBMF often provides more accurate inferences than other methods. In the GTEx data, EBMF identifies interpretable structure that agrees with known relationships among human tissues. Software implementing our approach is available at <https://github.com/stephenslab/flashr>.

Keywords: empirical Bayes, matrix factorization, normal means, sparse prior, unimodal prior, variational approximation

1. Introduction

Matrix factorization methods are widely used for inferring and summarizing structure in multivariate data. In brief, these methods represent an observed $n \times p$ data matrix Y as:

$$Y = L^T F + E \quad (1.1)$$

where L is an $K \times n$ matrix, F is a $K \times p$ matrix, and E is an $n \times p$ matrix of residuals (whose entries we assume to be normally distributed, although the methods we develop can be generalized to other settings; see Section 6.4). Here we adopt the notation and terminology of factor analysis, and refer to L as the “loadings” and F as the “factors”.

The model (1.1) has many potential applications. One range of applications arise from “matrix completion” problems (e.g. Fithian et al., 2018): methods that estimate L and F in (1.1) from partially observed Y provide a natural and effective way to fill in the missing entries. Another wide range of applications come from the desire to *summarize and understand* the structure in a matrix Y : in (1.1) each row of Y is approximated by a linear combination of underlying “factors” (rows of F), which—ideally—have some intuitive or scientific interpretation. For example, suppose Y_{ij} represents the rating of a user i for a movie j . Each factor might represent a genre of movie (“comedy”, “drama”, “romance”, “horror” etc), and the ratings for a user i could be written as a linear combination of these factors, with the weights (loadings) representing how much individual i likes that genre. Or, suppose Y_{ij} represents the expression of gene j in sample i . Each factor might represent a module of co-regulated genes, and the data for sample i could be written as a linear combination of these factors, with the loadings representing how active each module is in each sample. Many other examples could be given across many fields, including psychology (Ford et al., 1986), econometrics (Bai and Ng, 2008), natural language processing (Bouchard et al., 2015), population genetics (Engelhardt and Stephens, 2010), and functional genomics (Stein-O’Brien et al., 2018).

The simplest approaches to estimating L and/or F in (1.1) are based on maximum likelihood or least squares. For example, Principal Components Analysis (PCA)—or, more precisely, truncated Singular Value Decomposition (SVD)—can be interpreted as fitting (1.1) by least squares, assuming that columns of L are orthogonal and columns of F are orthonormal (Eckart and Young, 1936). And classical factor analysis (FA) corresponds to maximum likelihood estimation of L , assuming that the elements of F are independent standard normal and allowing different residual variances for each column of Y (Rubin and Thayer, 1982). While these simple methods remain widely used, in the last two decades researchers have focused considerable attention on obtaining more accurate and/or more interpretable estimates, either by imposing additional constraints (e.g. non-negativity; Lee and Seung, 1999) or by regularization using a penalty term (e.g. Jolliffe et al., 2003; Witten et al., 2009; Mazumder et al., 2010; Hastie et al., 2015; Fithian et al., 2018), or a prior distribution (e.g. Bishop, 1999; Attias, 1999; Ghahramani and Beal, 2000; West, 2003). In particular, many authors have noted the benefits of sparsity assumptions on L and/or F —particularly in applications where interpretability of the estimates is desired—and there now exists a wide range of methods that attempt to induce sparsity in these models (e.g. Sabatti and James, 2005; Zou et al., 2006; Pournara and Wernisch, 2007; Carvalho et al., 2008; Witten et al., 2009; Engelhardt and Stephens, 2010; Knowles and Ghahramani, 2011; Bhattacharya and Dunson, 2011; Mayrink et al., 2013; Yang et al., 2014; Gao et al., 2016; Hore et al., 2016; Ročková and George, 2016; Srivastava et al., 2017; Kaufmann and Schumacher, 2017; Frühwirth-Schnatter and Lopes, 2018; Zhao et al., 2018). Many of these methods induce sparsity in the loadings only, although some induce sparsity in both loadings and factors.

In any statistical problem involving sparsity, a key question is how strong the sparsity should be. In penalty-based methods this is controlled by the strength and form of the penalty, whereas in Bayesian methods it is controlled by the prior distributions. In this paper we take an Empirical Bayes approach to this problem, exploiting variational approximation methods (Blei et al., 2017) to obtain simple algorithms that jointly estimate

the prior distributions for both loadings and factors, as well as the loadings and factors themselves.

Both EB and variational methods have been previously used for this problem (Bishop, 1999; Lim and Teh, 2007; Raiko et al., 2007; Stegle et al., 2010). However, most of this previous work has used simple normal prior distributions that do not induce sparsity. Variational methods that use sparsity-inducing priors include Girolami (2001), which uses a Laplace prior on the factors (no prior on the loadings which are treated as free parameters), Hochreiter et al. (2010) which extends this to Laplace priors on both factors and loadings, with fixed values of the Laplace prior parameters; Titsias and Lázaro-Gredilla (2011), which uses a sparse “spike-and-slab” (point normal) prior on the loadings (with the same prior on all K loadings) and a normal prior on the factors; and Hore et al. (2016) which uses a spike-and-slab prior on one mode in a tensor decomposition. (While this work was in review further examples appeared, including Argelaguet et al. 2018, which uses normal priors on the loadings and point-normal on the factors.)

Our primary contribution here is to develop and implement a more general EB approach to matrix factorization (EBMF). This general approach allows for a wide range of potential sparsity-inducing prior distributions on both the loadings and the factors within a single algorithmic framework. We accomplish this by showing that, when using variational methods, fitting EBMF with *any* prior family can be reduced to repeatedly solving a much simpler problem—the “empirical Bayes normal means” (EBNM) problem—with the same prior family. This feature makes it easy to implement methods for any desired prior family—one simply has to implement a method to solve the corresponding normal means problem, and then plug this into our algorithm. This approach can work for both parametric families (e.g. normal, point-normal, laplace, point-laplace) and non-parametric families, including the “adaptive shrinkage” priors (unimodal and scale mixtures of normals) from Stephens (2017). It is also possible to accommodate non-negative constraints on either L and/or F by using non-negative prior families. Even simple versions of our approach—e.g. using point-normal priors on both factors and loadings—provide more generality than most existing EBMF approaches and software.

A second contribution of our work is to highlight similarities and differences between EBMF and penalty-based methods for regularizing L and/or F . Indeed, our algorithm for fitting EBMF has the same structure as commonly-used algorithms for penalty-based methods, with the prior distribution playing a role analogous to the penalty (see Remark 3 later). While the general correspondence between estimates from penalized methods and Bayesian posterior modes (MAP estimates) is well known, the connection here is different, because the EBMF approach is estimating a posterior mean, not a mode (indeed, with sparse priors the MAP estimates of L and F are not useful because they are trivially 0). A key difference between the EBMF approach and penalty-based methods is that the EBMF prior is estimated by solving an optimization problem, whereas in penalty-based methods the strength of the penalty is usually chosen by cross-validation. This difference makes it much easier for EBMF to allow for different levels of sparsity in every factor and every loading: in EBMF one simply uses a different prior for every factor and loading, whereas tuning a separate parameter for every factor and loading by CV becomes very cumbersome.

The final contribution is that we provide an R software package, *flash* (Factors and Loadings by Adaptive SHrinkage), implementing our flexible EBMF framework. We demonstrate

the utility of these methods through both numerical comparisons with competing methods and through a scientific application: analysis of data from the GTEx (Genotype Tissue Expression) project on genetic associations across 44 human tissues. In numerical comparisons *flash* often provides more accurate inferences than other methods, while remaining computationally tractable for moderate-sized matrices (millions of entries). In the GTEx data, *flash* highlight both effects that are shared across many tissues (“dense” factors) and effects that are specific to a small number of tissues (“sparse” factors). These sparse factors often highlight similarities between tissues that are known to be biologically related, providing external support for the reliability of the results.

2. A General Empirical Bayes Matrix Factorization Model

We define the K -factor Empirical Bayes Matrix Factorization (EBMF) model as follows:

$$Y = \sum_{k=1}^K \mathbf{l}_k \mathbf{f}_k^T + E \quad (2.1)$$

$$l_{k1}, \dots, l_{kn} \stackrel{iid}{\sim} g_{\mathbf{l}_k}, \quad g_{\mathbf{l}_k} \in \mathcal{G}_{\mathbf{l}} \quad (2.2)$$

$$f_{k1}, \dots, f_{kp} \stackrel{iid}{\sim} g_{\mathbf{f}_k}, \quad g_{\mathbf{f}_k} \in \mathcal{G}_{\mathbf{f}} \quad (2.3)$$

$$E_{ij} \sim N(0, 1/\tau_{ij}) \text{ with } \boldsymbol{\tau} := (\tau_{ij}) \in \mathcal{T}. \quad (2.4)$$

Here Y is the $n \times p$ observed data matrix, \mathbf{l}_k is an n -vector (the k th set of “loadings”), \mathbf{f}_k is a p -vector (the k th “factor”), $\mathcal{G}_{\mathbf{l}}$ and $\mathcal{G}_{\mathbf{f}}$ are pre-specified (possibly non-parametric) families of distributions, $g_{\mathbf{l}_k}$ and $g_{\mathbf{f}_k}$ are unknown “prior” distributions that are to be estimated, E is an $n \times p$ matrix of independent error terms, and $\boldsymbol{\tau}$ is an unknown $n \times p$ matrix of precisions (τ_{ij}) which is assumed to lie in some space \mathcal{T} . (This allows structure to be imposed on $\boldsymbol{\tau}$, such as constant precision, $\tau_{ij} = \tau$, or column-specific precisions, $\tau_{ij} = \tau_j$, for example.) Our methods allow that some elements of Y may be “missing”, and can estimate the missing values (Section 4.1).

The term “Empirical Bayes” in EBMF means we fit (2.1)-(2.4) by obtaining *point estimates* for the priors $g_{\mathbf{l}_k}, g_{\mathbf{f}_k}$, ($k = 1, \dots, K$) and approximate the posterior distributions for the parameters $\mathbf{l}_k, \mathbf{f}_k$ given those point estimates. This contrasts with a “fully Bayes” approach that, instead of obtaining point estimates for $g_{\mathbf{l}_k}, g_{\mathbf{f}_k}$, would integrate over uncertainty in the estimates. This would involve specifying prior distributions for $g_{\mathbf{l}_k}, g_{\mathbf{f}_k}$ as well as (perhaps substantial) additional computation. The EB approach has the advantage of simplicity – both conceptually and computationally—while enjoying many of the benefits of a fully Bayes approach. In particular it allows for sharing of information across elements of each loading/factor. For example, if the data suggest that a particular factor, \mathbf{f}_k , is sparse, then this will be reflected in a sparse estimate of $g_{\mathbf{f}_k}$, and subsequently strong shrinkage of the smaller elements of f_{k1}, \dots, f_{kp} towards 0. Conversely, when the data suggest a non-sparse factor then the prior will be dense and the shrinkage less strong. By allowing different prior distributions for each factor and each loading, the model has the flexibility to adapt to any combination of sparse and dense loadings and factors. However, to fully capitalize on this flexibility one needs suitably flexible prior families $\mathcal{G}_{\mathbf{l}}$ and $\mathcal{G}_{\mathbf{f}}$ capable of capturing both sparse and dense factors. A key feature of our work is it allows for very flexible prior families, including non-parametric families.

Some specific choices of the distributional families \mathcal{G}_l and \mathcal{G}_f correspond to models used in previous work. In particular, many previous papers have studied the case with normal priors, where \mathcal{G}_l and \mathcal{G}_f are both the family of zero-mean normal distributions (e.g. Bishop, 1999; Lim and Teh, 2007; Raiko et al., 2007; Nakajima and Sugiyama, 2011). This family is particularly simple, having a single hyper-parameter, the prior variance, to estimate for each factor. However, it does not induce sparsity on either L or F ; indeed, when the matrix Y is fully observed, the estimates of L and F under a normal prior (when using a fully factored variational approximation) are simply scalings of the singular vectors from an SVD of Y (Nakajima and Sugiyama, 2011; Nakajima et al., 2013). Our work here extends these previous approaches to a much wider range of prior families that *do* induce sparsity on L and/or F .

We note that the EBMF model (2.1)-(2.4) differs in an important way from the sparse factor analysis (SFA) methods in Engelhardt and Stephens (2010), which use a type of Automatic Relevance Determination prior (e.g. Tipping, 2001; Wipf and Nagarajan, 2008) to induce sparsity on the loadings matrix. In particular, SFA estimates a separate hyper-parameter for every element of the loadings matrix, with no sharing of information across elements of the same loading. In contrast, EBMF estimates a single *shared* prior distribution for elements of each loading, which, as noted above, allows for sharing of information across elements of each loading/factor.

3. Fitting the EBMF Model

To simplify exposition we begin with the case $K = 1$ (“rank 1”); see Section 3.6 for the extension to general K . To simplify notation we assume the families $\mathcal{G}_l, \mathcal{G}_f$ are the same, so we can write $\mathcal{G}_l = \mathcal{G}_f = \mathcal{G}$. To further lighten notation in the case $K = 1$ we use $g_l, g_f, \mathbf{l}, \mathbf{f}$ instead of $g_{l_1}, g_{f_1}, \mathbf{l}_1, \mathbf{f}_1$.

Fitting the EBMF model involves estimating all of $g_l, g_f, \mathbf{l}, \mathbf{f}, \boldsymbol{\tau}$. A standard EB approach would be to do this in two steps:

- Estimate g_l, g_f and $\boldsymbol{\tau}$, by maximizing the likelihood:

$$L(g_l, g_f, \boldsymbol{\tau}) := \int \int p(Y|\mathbf{l}, \mathbf{f}, \boldsymbol{\tau}) g_l(d\mathbf{l}_1) \dots g_l(d\mathbf{l}_n) g_f(d\mathbf{f}_1) \dots g_f(d\mathbf{f}_p) \quad (3.1)$$

over $g_l, g_f \in \mathcal{G}, \boldsymbol{\tau} \in \mathcal{T}$. (This optimum will typically not be unique because of identifiability issues; see Section 3.8.)

- Estimate \mathbf{l} and \mathbf{f} using their posterior distribution: $p(\mathbf{l}, \mathbf{f}|Y, \hat{g}_l, \hat{g}_f, \hat{\boldsymbol{\tau}})$.

However, both these two steps are difficult, even for very simple choices of \mathcal{G} . Instead, following previous work (see Introduction for citations) we use variational approximations to approximate this approach. Although variational approximations are known to typically under-estimate uncertainty in posterior distributions, our focus here is on obtaining useful point estimates for \mathbf{l}, \mathbf{f} ; results shown later demonstrate that the variational approximation can perform well in this task.

3.1 The Variational Approximation

The variational approach—see Blei et al. (2017) for review—begins by writing the log of the likelihood (3.1) as:

$$l(g_l, g_f, \boldsymbol{\tau}) := \log L(g_l, g_f, \boldsymbol{\tau}) \tag{3.2}$$

$$= F(q, g_l, g_f, \boldsymbol{\tau}) + D_{KL}(q||p) \tag{3.3}$$

where

$$F(q, g_l, g_f, \boldsymbol{\tau}) = \int q(\mathbf{l}, \mathbf{f}) \log \frac{p(Y, \mathbf{l}, \mathbf{f} | g_l, g_f, \boldsymbol{\tau})}{q(\mathbf{l}, \mathbf{f})} d\mathbf{l} d\mathbf{f}, \tag{3.4}$$

and

$$D_{KL}(q||p) = - \int q(\mathbf{l}, \mathbf{f}) \log \frac{p(\mathbf{l}, \mathbf{f} | Y, g_l, g_f, \boldsymbol{\tau})}{q(\mathbf{l}, \mathbf{f})} d\mathbf{l} d\mathbf{f} \tag{3.5}$$

is the Kullback–Leibler divergence from q to p . This identity holds for any distribution $q(\mathbf{l}, \mathbf{f})$. Because D_{KL} is non-negative, it follows that $F(q, g_l, g_f, \boldsymbol{\tau})$ is a lower bound for the log likelihood:

$$l(g_l, g_f, \boldsymbol{\tau}) \geq F(q, g_l, g_f, \boldsymbol{\tau}) \tag{3.6}$$

with equality when $q(\mathbf{l}, \mathbf{f}) = p(\mathbf{l}, \mathbf{f} | Y, g_l, g_f, \boldsymbol{\tau})$.

In other words,

$$l(g_l, g_f, \boldsymbol{\tau}) = \max_q F(q, g_l, g_f, \boldsymbol{\tau}), \tag{3.7}$$

where the maximization is over all possible distributions $q(\mathbf{l}, \mathbf{f})$. Maximizing $l(g_l, g_f, \boldsymbol{\tau})$ can thus be viewed as maximizing F over $q, g_l, g_f, \boldsymbol{\tau}$. However, as noted above, this maximization is difficult. The variational approach simplifies the problem by maximizing F but restricting the family of distributions for q . Specifically, the most common variational approach—and the one we consider here—restricts q to the family \mathcal{Q} of distributions that “fully-factorize”:

$$\mathcal{Q} = \left\{ q : q(\mathbf{l}, \mathbf{f}) = \prod_{i=1}^n q_{l,i}(l_i) \prod_{j=1}^p q_{f,j}(f_j) \right\}. \tag{3.8}$$

The variational approach seeks to optimize F over $q, g_l, g_f, \boldsymbol{\tau}$ with the constraint $q \in \mathcal{Q}$. For $q \in \mathcal{Q}$ we can write $q(\mathbf{l}, \mathbf{f}) = q_l(\mathbf{l})q_f(\mathbf{f})$ where $q_l(\mathbf{l}) = \prod_{i=1}^n q_{l,i}(l_i)$ and $q_f(\mathbf{f}) = \prod_{j=1}^p q_{f,j}(f_j)$, and we can consider the problem as maximizing $F(q_l, q_f, g_l, g_f, \boldsymbol{\tau})$.

3.2 Alternating Optimization

We optimize $F(q_l, q_f, g_l, g_f, \boldsymbol{\tau})$ by alternating between optimizing over variables related to \mathbf{l} [(q_l, g_l)], over variables related to \mathbf{f} [(q_f, g_f)], and over $\boldsymbol{\tau}$. Each of these steps is guaranteed to increase (or, more precisely, not decrease) F , and convergence can be assessed by (for example) stopping when these optimization steps yield a very small increase in F . Note that F may be multi-modal, and there is no guarantee that the algorithm will converge to a global optimum. The approach is summarized in Algorithm 1.

Algorithm 1 Alternating Optimization for EBMF (rank 1)

Require: Initial values $q_l^{(0)}, q_f^{(0)}, g_l^{(0)}, g_f^{(0)}$
 1: $t \leftarrow 0$
 2: **repeat**
 3: $t \leftarrow t + 1$
 4: $\boldsymbol{\tau}^{(t)} \leftarrow \arg \max_{\boldsymbol{\tau}} F(q_l^{(t-1)}, q_f^{(t-1)}, g_l^{(t-1)}, g_f^{(t-1)}, \boldsymbol{\tau})$
 5: $q_l^{(t)}, g_l^{(t)} \leftarrow \arg \max_{q_l, g_l} F(q_l, q_f^{(t-1)}, g_l, g_f^{(t-1)}, \boldsymbol{\tau}^{(t)})$.
 6: $q_f^{(t)}, g_f^{(t)} \leftarrow \arg \max_{q_f, g_f} F(q_l^{(t)}, q_f, g_l^{(t)}, g_f, \boldsymbol{\tau}^{(t)})$.
 7: **until** converged
 8: **return** $q_l^{(t)}, q_f^{(t)}, g_l^{(t)}, g_f^{(t)}, \boldsymbol{\tau}^{(t)}$

The key steps in Algorithm 1 are the maximizations in Steps 4-6.

Step 4, the update of $\boldsymbol{\tau}$, involves computing the expected squared residuals:

$$\overline{R}_{ij}^2 := \mathbb{E}_{q_l, q_f} [(Y_{ij} - l_i f_j)^2] \quad (3.9)$$

$$= [Y_{ij} - \mathbb{E}_{q_l}(l_i) \mathbb{E}_{q_f}(f_j)]^2 - \mathbb{E}_{q_l}(l_i)^2 \mathbb{E}_{q_f}(f_j)^2 + \mathbb{E}_{q_l}(l_i^2) \mathbb{E}_{q_f}(f_j^2). \quad (3.10)$$

This is straightforward provided the first and second moments of q_l and q_f are available (see Appendix A.1 for details).

Steps 5 and 6 are essentially identical except for switching the role of l and f . One of our key results is that each of these steps can be achieved by solving a simpler problem—the Empirical Bayes normal means (EBNM) problem. The next subsection (3.3) describes the EBNM problem, and the following subsection (3.4) details how this can be used to solve Steps 5 and 6.

3.3 The EBNM Problem

Suppose we have observations $\mathbf{x} = (x_1, \dots, x_n)$ of underlying quantities $\boldsymbol{\theta} = (\theta_1, \dots, \theta_n)$, with independent Gaussian errors with known standard deviations $\mathbf{s} = (s_1, \dots, s_n)$. Suppose further that the elements of $\boldsymbol{\theta}$ are assumed i.i.d. from some distribution, $g \in \mathcal{G}$. That is,

$$\mathbf{x} | \boldsymbol{\theta} \sim N_n(\boldsymbol{\theta}, \text{diag}(s_1^2, \dots, s_n^2)) \quad (3.11)$$

$$\theta_1, \dots, \theta_n \sim^{iid} g, \quad g \in \mathcal{G}, \quad (3.12)$$

where $N_n(\boldsymbol{\mu}, \Sigma)$ denotes the n -dimensional normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix Σ .

By solving the EBNM problem we mean fitting the model (3.11)-(3.12) by the following two-step procedure:

1. Estimate g by maximum (marginal) likelihood:

$$\hat{g} = \arg \max_{g \in \mathcal{G}} \prod_j \int p(x_j | \theta_j, s_j) g(d\theta_j). \quad (3.13)$$

2. Compute the posterior distribution for $\boldsymbol{\theta}$ given \hat{g} ,

$$p(\boldsymbol{\theta}|\mathbf{x}, \mathbf{s}, \hat{g}) \propto \prod_j \hat{g}(\theta_j)p(x_j|\theta_j, s_j). \quad (3.14)$$

Later in this paper we will have need for the posterior first and second moments, so we define them here for convenience:

$$\bar{\theta}_j := E(\theta_j|\mathbf{x}, \mathbf{s}, \hat{g}) \quad (3.15)$$

$$\bar{\theta}_j^2 := E(\theta_j^2|\mathbf{x}, \mathbf{s}, \hat{g}). \quad (3.16)$$

Formally, this procedure defines a mapping (which depends on the family \mathcal{G}) from the known quantities (\mathbf{x}, \mathbf{s}) , to (\hat{g}, p) , where \hat{g}, p are given in (3.13) and (3.14). We use *EBNM* to denote this mapping:

$$EBNM(\mathbf{x}, \mathbf{s}) = (\hat{g}, p). \quad (3.17)$$

Remark 1 *Solving the EBNM problem is central to all our algorithms, so it is worth some study. A key point is that the EBNM problem provides an attractive and flexible way to induce shrinkage and/or sparsity in estimates of θ . For example, if θ is truly sparse, with many elements at or near 0, then the estimate \hat{g} will typically have considerable mass near 0, and the posterior means (3.15) will be “shrunk” strongly toward 0 compared with the original observations. In this sense solving the EBNM problem can be thought of as a model-based analogue of thresholding-based methods, with the advantage that by estimating g from the data the EBNM approach automatically adapts to provide an appropriate level of shrinkage. These ideas have been used in wavelet denoising (Clyde and George, 2000; Johnstone et al., 2004; Johnstone and Silverman, 2005a; Xing et al., 2016), and false discovery rate estimation (Thomas et al., 1985; Stephens, 2017) for example. Here we apply them to matrix factorization problems.*

3.4 Connecting the EBMF and EBNM Problems

The EBNM problem is well studied, and can be solved reasonably easily for many choices of \mathcal{G} (e.g. Johnstone and Silverman, 2005b; Koenker and Mizera, 2014a; Stephens, 2017). In Section 4 we give specific examples; for now our main point is that *if* one can solve the EBNM problem for a particular choice of \mathcal{G} then it can be used to implement Steps 5 and 6 in Algorithm 1 for the corresponding EBMF problem. The following Proposition formalizes this for Step 5 of Algorithm 1; a similar proposition holds for Step 6 (see also Appendix A).

Proposition 2 *Step 5 in Algorithm 1 is solved by solving an EBNM problem. Specifically*

$$\arg \max_{q_l, q_f, g_l, g_f, \boldsymbol{\tau}} F(q_l, q_f, g_l, g_f, \boldsymbol{\tau}) = EBNM(\hat{l}(Y, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, \boldsymbol{\tau}), \mathbf{s}_l(\bar{\mathbf{f}}^2, \boldsymbol{\tau})) \quad (3.18)$$

where the functions $\hat{\mathbf{l}} : R^{n \times p} \times R^p \times R^p \times R^{n \times p} \rightarrow R^n$ and $\mathbf{s}_l : R^p \times R^{n \times p} \rightarrow R^n$ are given by

$$\hat{\mathbf{l}}(Y, \mathbf{v}, \mathbf{w}, \boldsymbol{\tau})_i := \frac{\sum_j \tau_{ij} Y_{ij} v_j}{\sum_j \tau_{ij} w_j}, \quad (3.19)$$

$$\mathbf{s}_l(\mathbf{w}, \boldsymbol{\tau})_i := \left(\sum_j \tau_{ij} w_j \right)^{-0.5}, \quad (3.20)$$

and $\bar{\mathbf{f}}, \bar{\mathbf{f}}^2 \in R^p$ denote the vectors whose elements are the first and second moments of \mathbf{f} under $q_{\mathbf{f}}$:

$$\bar{\mathbf{f}} := (E_{q_{\mathbf{f}}}(f_j)) \quad (3.21)$$

$$\bar{\mathbf{f}}^2 := (E_{q_{\mathbf{f}}}(f_j^2)). \quad (3.22)$$

Proof See Appendix A. ■

For intuition into where the EBNM in Proposition 2 comes from, consider estimating \mathbf{l}, g_l in (2.1) with \mathbf{f} and $\boldsymbol{\tau}$ known. The model then becomes n independent regressions of the rows of Y on \mathbf{f} , and the maximum likelihood estimate for \mathbf{l} has elements:

$$\hat{l}_i = \frac{\sum_j \tau_{ij} Y_{ij} f_j}{\sum_j \tau_{ij} f_j^2}, \quad (3.23)$$

with standard errors

$$s_i = \left(\sum_j \tau_{ij} f_j^2 \right)^{-0.5}. \quad (3.24)$$

Further, it is easy to show that

$$\hat{l}_i \sim N(l_i, s_i^2). \quad (3.25)$$

Combining (3.25) with the prior

$$l_1, \dots, l_n \sim^{iid} g_l, \quad g_l \in \mathcal{G} \quad (3.26)$$

yields an EBNM problem.

The EBNM in Proposition 2 is the same as the EBNM (3.25)-(3.26), but with the terms f_j and f_j^2 replaced with their expectations under $q_{\mathbf{f}}$. Thus, the update for (q_l, g_l) in Algorithm 1, with $(q_{\mathbf{f}}, g_{\mathbf{f}}, \boldsymbol{\tau})$ fixed, is closely connected to solving the EBMF problem for “known $\mathbf{f}, \boldsymbol{\tau}$ ”.

3.5 Streamlined Implementation Using First and Second Moments

Although Algorithm 1, as written, optimizes over $(q_l, q_{\mathbf{f}}, g_l, g_{\mathbf{f}})$, in practice each step requires only the first and second moments of the distributions q_l and $q_{\mathbf{f}}$. For example, the EBNM problem in Proposition 1 involves $\bar{\mathbf{f}}$ and $\bar{\mathbf{f}}^2$ and not $g_{\mathbf{f}}$. Consequently, we can

simplify implementation by keeping track of only those moments. In particular, when solving the normal means problem, $EBNM(\mathbf{x}, \mathbf{s})$ in (3.17), we need only return the posterior first and second moments (3.15) and (3.16). This results in a streamlined and intuitive implementation, summarized in Algorithm 2.

Algorithm 2 Streamlined Alternating Optimization for EBMF (rank 1)

Require: A data matrix Y ($n \times p$)

Require: A function, $\text{ebnm}(\mathbf{x}, \mathbf{s}) \rightarrow (\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}}^2)$, that solves the EBNM problem (3.11)-(3.12) and returns the first and second posterior moments (3.15)-(3.16).

Require: A function, $\text{init}(Y) \rightarrow (\hat{\mathbf{l}}, \hat{\mathbf{f}})$ that produces initial estimates for \mathbf{l} (an n vector) and \mathbf{f} (a p vector) given data Y . (For example, rank 1 singular value decomposition.)

- 1: Initialize first moments $(\bar{\mathbf{l}}, \bar{\mathbf{f}})$, using $(\bar{\mathbf{l}}, \bar{\mathbf{f}}) \leftarrow \text{init}(Y)$
 - 2: Initialize second moments $(\bar{\mathbf{l}}^2, \bar{\mathbf{f}}^2)$, by squaring first moments: $\bar{\mathbf{l}}^2 \leftarrow (\bar{\mathbf{l}}_i^2)$ and $\bar{\mathbf{f}}^2 \leftarrow (\bar{\mathbf{f}}_j^2)$.
 - 3: **repeat**
 - 4: Compute the matrix of expected squared residuals \bar{R}_{ij}^2 from (3.9).
 - 5: $\tau_j \leftarrow n / \sum_i \bar{R}_{ij}^2$. [This update assumes column-specific variances; it can be modified to make other assumptions.]
 - 6: Compute $\hat{\mathbf{l}}(Y, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, \boldsymbol{\tau})$ and standard errors $\mathbf{s}_l(\bar{\mathbf{l}}^2, \boldsymbol{\tau})$, using (3.19) and (3.20).
 - 7: $(\bar{\mathbf{l}}, \bar{\mathbf{l}}^2) \leftarrow \text{ebnm}(\hat{\mathbf{l}}, \mathbf{s}_l)$.
 - 8: Compute $\hat{\mathbf{f}}(Y, \bar{\mathbf{l}}, \bar{\mathbf{l}}^2, \boldsymbol{\tau})$ and standard errors $\mathbf{s}_f(\bar{\mathbf{l}}^2, \boldsymbol{\tau})$ (similarly as for $\hat{\mathbf{l}}$ and \mathbf{s}_l ; see (A.14) and (A.15)).
 - 9: $(\bar{\mathbf{f}}, \bar{\mathbf{f}}^2) \leftarrow \text{ebnm}(\hat{\mathbf{f}}, \mathbf{s}_f)$.
 - 10: **until** converged
 - 11: **return** $\bar{\mathbf{l}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, \boldsymbol{\tau}$
-

Remark 3 *Algorithm 2 has a very intuitive form: it has the flavor of an alternating least squares algorithm, which alternates between estimating \mathbf{l} given \mathbf{f} (Step 6) and \mathbf{f} given \mathbf{l} (Step 8), but with the addition of the ebnm step (Steps 7 and 9), which can be thought of as regularizing or shrinking the estimates: see Remark 1. This viewpoint highlights connections with related algorithms. For example, the (rank 1 version of the) SSVD algorithm from Yang et al. (2014) has a similar form, but uses a thresholding function in place of the ebnm function to induce shrinkage and/or sparsity.*

3.6 The K -factor EBMF Model

It is straightforward to extend the variational approach to fit the general K factor model (2.1)-(2.4). In brief, we introduce variational distributions $(q_{\mathbf{l}_k}, q_{\mathbf{f}_k})$ for $k = 1, \dots, K$, and then optimize the objective function $F(q_{\mathbf{l}_1}, g_{\mathbf{l}_1}, q_{\mathbf{f}_1}, g_{\mathbf{f}_1}; \dots; q_{\mathbf{l}_K}, g_{\mathbf{l}_K}, q_{\mathbf{f}_K}, g_{\mathbf{f}_K}; \boldsymbol{\tau})$. Similar to the rank-1 model, this optimization can be done by iteratively updating parameters relating to a single loading or factor, keeping other parameters fixed. And again we simplify implementation by keeping track of only the first and second moments of the distributions $q_{\mathbf{l}_k}$ and $q_{\mathbf{f}_k}$, which we denote $\bar{\mathbf{l}}_k, \bar{\mathbf{l}}_k^2, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_k^2$. The updates to $\bar{\mathbf{l}}_k, \bar{\mathbf{l}}_k^2$ (and $\bar{\mathbf{f}}_k, \bar{\mathbf{f}}_k^2$) are essentially identical to those for fitting the rank 1 model above, but with Y_{ij} replaced with

the residuals obtained by removing the estimated effects of the other $k - 1$ factors:

$$R_{ij}^k := Y_{ij} - \sum_{k' \neq k} \bar{u}_{k'i} \bar{f}_{k'j}. \quad (3.27)$$

Based on this approach we have implemented two algorithms for fitting the K -factor model. First, a simple “greedy” algorithm, which starts by fitting the rank 1 model, and then adds factors $k = 2, \dots, K$, one at a time, optimizing over the new factor parameters before moving on to the next factor. Second, a “backfitting” algorithm (Breiman and Friedman, 1985), which iteratively refines the estimates for each factor given the estimates for the other factors. Both algorithms are detailed in Appendix A.

3.7 Selecting K

An interesting feature of EB approaches to matrix factorization, noted by Bishop (1999), is that they automatically select the number of factors K . This is because the maximum likelihood solution to $g_{\mathbf{u}_k}, g_{\mathbf{f}_k}$ is sometimes a point mass on 0 (provided \mathcal{G} includes this distribution). Furthermore, the same is true of the solution to the variational approximation (see also Bishop, 1999; Stegle et al., 2012). This means that if K is set sufficiently large then some loading/factor combinations will be optimized to be exactly 0. (Or, in the greedy approach, which adds one factor at a time, the algorithm will eventually add a factor that is exactly 0, at which point it terminates.)

Here we note that the variational approximation may be expected to result in conservative estimation (i.e. underestimation) of K compared with the (intractable) use of maximum likelihood to estimate $g_{\mathbf{u}}, g_{\mathbf{f}}$. We base our argument on the simplest case: comparing $K = 1$ vs $K = 0$. Let δ_0 denote the degenerate distribution with all its mass at 0. Note that the rank-1 factor model (2.1), with $g_{\mathbf{u}} = \delta_0$ (or $g_{\mathbf{f}} = \delta_0$) is essentially a “rank-0” model. Now note that the variational lower bound, F , is exactly equal to the log-likelihood when $g_{\mathbf{u}} = \delta_0$ (or $g_{\mathbf{f}} = \delta_0$). This is because if the prior is a point mass at 0 then the posterior is also a point mass, which trivially factorizes as a product of point masses, and so the variational family \mathcal{Q} includes the true posterior in this case. Since F is a lower bound to the log-likelihood we have the following simple lemma:

Lemma 4 *If $F(\hat{q}, \hat{g}_{\mathbf{u}}, \hat{g}_{\mathbf{f}}, \hat{\tau}) > F(\delta_0, \delta_0, \delta_0, \hat{\tau}_0)$ then $l(\hat{g}_{\mathbf{u}}, \hat{g}_{\mathbf{f}}, \hat{\tau}) > l(\delta_0, \delta_0, \hat{\tau}_0)$.*

Proof

$$l(\hat{g}_{\mathbf{u}}, \hat{g}_{\mathbf{f}}, \hat{\tau}) \geq F(\hat{q}, \hat{g}_{\mathbf{u}}, \hat{g}_{\mathbf{f}}, \hat{\tau}) > F(\delta_0, \delta_0, \delta_0, \hat{\tau}_0) = l(\delta_0, \delta_0, \hat{\tau}_0) \quad (3.28)$$

■

Thus, if the variational approximation F favors $\hat{g}_{\mathbf{u}}, \hat{g}_{\mathbf{f}}, \hat{\tau}$ over the rank 0 model, then it is guaranteed that the likelihood would also favor $\hat{g}_{\mathbf{u}}, \hat{g}_{\mathbf{f}}, \hat{\tau}$ over the rank 0 model. In other words, compared with the likelihood, the variational approximation is conservative in terms of preferring the rank 1 model to the rank 0 model. This conservatism is a double-edged sword. On the one hand it means that if the variational approximation finds structure it should be taken seriously. On the other hand it means that the variational approximation could miss subtle structure.

In practice Algorithm 2 can converge to a local optimum of F that is not as high as the trivial (rank 0) solution, $F(\delta_0, \delta_0, \delta_0, \hat{\tau}_0)$. We can add a check for this at the end of Algorithm 2, and set $\hat{g}_l = \hat{g}_f = \delta_0$ and $\hat{\tau} = \hat{\tau}_0$ when this occurs.

3.8 Identifiability

In EBMF each loading and factor is identifiable, at best, only up to a multiplicative constant (provided \mathcal{G} is a scale family). Specifically, scaling the prior distributions g_{f_k} and g_{l_k} by c_k and $1/c_k$ respectively results in the same marginal likelihood, and also results in a corresponding scaling of the posterior distribution on the factors f_k and loadings l_k (e.g. it scales the posterior first moments by $c_k, 1/c_k$ and the second moments by $c_k^2, 1/c_k^2$). However, this non-identifiability is not generally a problem, and if necessary it could be dealt with by re-scaling factor estimates to have norm 1.

4. Software Implementation: *flash*

We have implemented Algorithms 2, 4 and 5 in an R package, *flash* (“factors and loadings via adaptive shrinkage”). These algorithms can fit the EBMF model for any choice of distributional family $\mathcal{G}_l, \mathcal{G}_f$: the user must simply provide a function to solve the EBNM problem for these prior families.

One source of functions for solving the EBNM problem is the “adaptive shrinkage” (*ashr*) package, which implements methods from Stephens (2017). These methods solve the EBNM problem for several flexible choices of \mathcal{G} , including:

- $\mathcal{G} = \mathcal{SN}$, the set of all scale mixtures of zero-centered normals;
- $\mathcal{G} = \mathcal{SU}$, the set of all symmetric unimodal distributions, with mode at 0;
- $\mathcal{G} = \mathcal{U}$, the set of all unimodal distributions, with mode at 0;
- $\mathcal{G} = \mathcal{U}_+$, the set of all non-negative unimodal distributions, with mode at 0.

These methods are computationally stable and efficient, being based on convex optimization methods (Koenker and Mizera, 2014b) and analytic Bayesian posterior computations.

We have also implemented functions to solve the EBNM problem for additional choices of \mathcal{G} in the package *ebnm* (<https://github.com/stephenslab/ebnm>). These include \mathcal{G} being the “point-normal” family:

- $\mathcal{G} = \mathcal{PN}$, the set of all distributions that are a mixture of a point mass at zero and a normal with mean 0.

This choice is less flexible than those in *ashr*, and involves non-convex optimizations, but can be faster.

Although in this paper we focus our examples on sparsity-inducing priors with $\mathcal{G}_l = \mathcal{G}_f = \mathcal{G}$ we note that our software makes it easy to experiment with different choices, some of which represent novel methodologies. For example, setting $\mathcal{G}_l = \mathcal{U}_+$ and $\mathcal{G}_f = \mathcal{SN}$ yields an EB version of semi-non-negative matrix factorization (Ding et al., 2008), and we are aware of no existing EB implementations for this problem. Exploring the relative merits of these many possible options in different types of application will be an interesting direction for future work.

4.1 Missing Data

If some elements of Y are missing, then this is easily dealt with. For example, the sums over j in (3.19) and (3.20) are simply computed using only the j for which Y_{ij} is not missing. This corresponds to an assumption that the missing elements of Y are “missing at random” (Rubin, 1976). In practice we implement this by setting $\tau_{ij} = 0$ whenever Y_{ij} is missing (and filling in the missing entries of Y to an arbitrary number). This allows the implementation to exploit standard fast matrix multiplication routines, which cannot handle missing data. If many data points are missing then it may be helpful to exploit sparse matrix routines.

4.2 Initialization

Both Algorithms 2 and 4 require a rank 1 initialization procedure, `init`. Here, we use the `softImpute` function from the package `softImpute` (Mazumder et al., 2010), with penalty parameter $\lambda = 0$, which essentially performs SVD when Y is completely observed, but can also deal with missing values in Y .

The backfitting algorithm (Algorithm 5) also requires initialization. One option is to use the greedy algorithm to initialize, which we call “greedy+backfitting”.

5. Numerical Comparisons

We now compare our methods with several competing approaches. To keep these comparisons manageable in scope we focus attention on methods that aim to capture possible sparsity in L and/or F . For EBMF we present results for two different shrinkage-oriented prior families, \mathcal{G} : the scale mixture of normals ($\mathcal{G} = \mathcal{SN}$), and the point-normal family ($\mathcal{G} = \mathcal{PN}$). We denote these *flash* and *flash_pn* respectively when we need to distinguish. In addition we consider Sparse Factor Analysis (SFA) (Engelhardt and Stephens, 2010), SFAMix (Gao et al., 2013), Nonparametric Bayesian Sparse Factor Analysis (NBSFA) (Knowles and Ghahramani, 2011), Penalized Matrix Decomposition (Witten et al., 2009) (PMD, implemented in the R package `PMA`), and Sparse SVD (Yang et al., 2014) (SSVD, implemented in R package `ssvd`). Although the methods we compare against involve only a small fraction of the very large number of methods for this problem, the methods were chosen to represent a wide range of different approaches to inducing sparsity: SFA, SFAMix and NBSFA are three Bayesian approaches with quite different approaches to prior specification; PMD is based on a penalized likelihood with L_1 penalty on factors and/or loadings; and SSVD is based on iterative thresholding of singular vectors. We also compare with `softImpute` (Mazumder et al., 2010), which does not explicitly model sparsity in L and F , but fits a regularized low-rank matrix using a nuclear-norm penalty. Finally, for reference we also use standard (truncated) SVD.

All of the Bayesian methods (*flash*, SFA, SFAMix and NBSFA) are “self-tuning”, at least to some extent, and we applied them here with default values. According to Yang et al. (2014) SSVD is robust to choice of tuning parameters, so we also ran SSVD with its default values, using the robust option (`method="method"`). The `softImpute` method has a single tuning parameter (λ , which controls the nuclear norm penalty), and we chose this penalty by orthogonal cross-validation (OCV; Appendix B). The PMD method can use two tuning parameters (one for \mathbf{l} and one for \mathbf{f}) to allow different sparsity levels in \mathbf{l} vs \mathbf{f} .

However, since tuning two parameters can be inconvenient it also has the option to use a single parameter for both \mathbf{l} and \mathbf{f} . We used OCV to tune parameters in both cases, referring to the methods as PMD.cv2 (2 tuning parameters) and PMD.cv1 (1 tuning parameter).

5.1 Simple Simulations

5.1.1 A SINGLE FACTOR EXAMPLE

We simulated data with $n = 200, p = 300$ under the single-factor model (2.1) with sparse loadings, and a non-sparse factor:

$$l_i \sim \pi_0 \delta_0 + (1 - \pi_0) \sum_{m=1}^5 \frac{1}{5} N(0, \sigma_m^2) \quad (5.1)$$

$$f_j \sim N(0, 1) \quad (5.2)$$

where δ_0 denotes a point mass on 0, and $(\sigma_1^2, \dots, \sigma_5^2) := (0.25, 0.5, 1, 2, 4)$. We simulated using three different levels of sparsity on the loadings, using $\pi_0 = 0.9, 0.3, 0$. (We set the noise precision $\tau = 1, 1/16, 1/25$ in these three cases to make each problem not too easy and not too hard.)

We applied all methods to this rank-1 problem, specifying the true value $K = 1$. (The NBSFA software does not provide the option to fix K , so is omitted here.) We compare methods in their accuracy in estimating the true low-rank structure ($B := \mathbf{l}\mathbf{f}^T$) using relative root mean squared error:

$$\text{RRMSE}(\hat{B}, B) := \sqrt{\frac{\sum_{i,j} (\hat{B}_{ij} - B_{ij})^2}{\sum_{i,j} B_{ij}^2}}. \quad (5.3)$$

Despite the simplicity of this simulation, the methods vary greatly in performance (Figure 1). Both versions of *flash* consistently outperform all the other methods across all scenarios (although *softImpute* performs similarly in the non-sparse case). The next best performances come from *softImpute* (SI.cv), PMD.cv2 and SFA, whose relative performances depend on the scenario. All three consistently improve on, or do no worse than, SVD. PMD.cv1 performs similarly to SVD. The SFAmix method performs very variably, sometimes providing very poor estimates, possibly due to poor convergence of the MCMC algorithm (it is the only method here that uses MCMC). The SSVD method consistently performs worse than simple SVD, possibly because it is more adapted to both factors and loadings being sparse (and possibly because, following Yang et al. 2014, we did not use CV to tune its parameters). Inspection of individual results suggests that the poor performance of both SFAmix and SSVD is often due to over-shrinking of non-zero loadings to zero.

5.1.2 A SPARSE BI-CLUSTER EXAMPLE (RANK 3)

An important feature of our EBMF methods is that they estimate separate distributions $g_{\mathbf{l}}, g_{\mathbf{f}}$ for each factor and each loading, allowing them to adapt to any combination of sparsity in the factors and loadings. This flexibility is not easy to achieve in other ways. For example, methods that use CV are generally limited to one or two tuning parameters because of the computational difficulties of searching over a larger space.

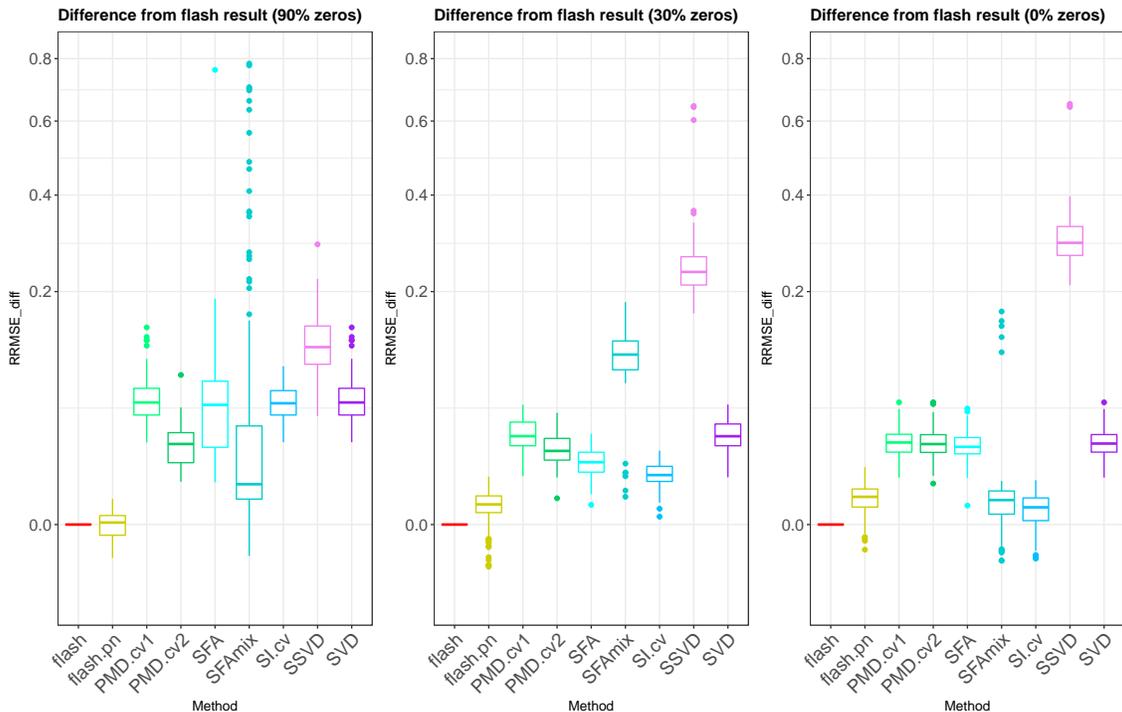


Figure 1: Boxplots comparing accuracy of *flash* with several other methods in a simple rank-1 simulation. This simulation involves a single dense factor, and a loading that varies from strong sparsity (90% zeros, left) to no sparsity (right). Accuracy is measured by difference in each methods RRMSE from the *flash* RRMSE, with smaller values indicating highest accuracy. The y axis is plotted on a non-linear (square-root) scale to avoid the plots being dominated by poorer-performing methods.

To illustrate this flexibility we simulated data under the factor model (2.1) with $n = 150$, $p = 240$, $K = 3$, $\tau = 1/4$, and:

$$\mathbf{l}_{1,i} \sim N(0, 2^2) \quad i = 1, \dots, 10 \quad (5.4)$$

$$\mathbf{l}_{2,i} \sim N(0, 1) \quad i = 11, \dots, 60 \quad (5.5)$$

$$\mathbf{l}_{3,i} \sim N(0, 1/2^2) \quad i = 61, \dots, 150 \quad (5.6)$$

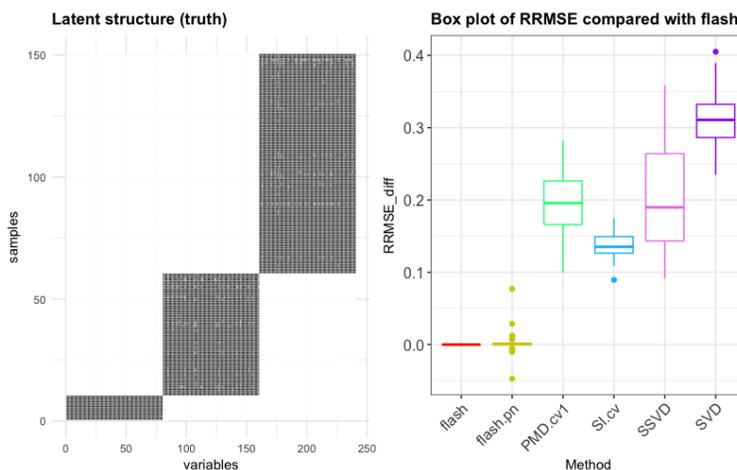
$$\mathbf{f}_{1,j} \sim N(0, 1/2^2) \quad j = 1, \dots, 80 \quad (5.7)$$

$$\mathbf{f}_{2,j} \sim N(0, 1) \quad j = 81, \dots, 160 \quad (5.8)$$

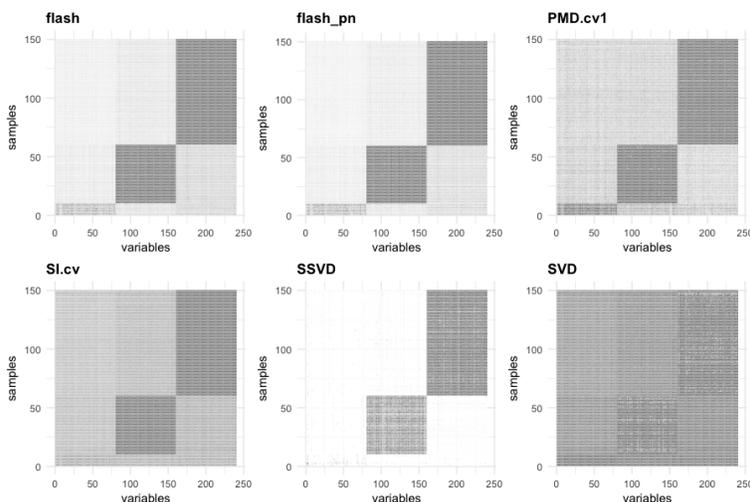
$$\mathbf{f}_{3,j} \sim N(0, 2^2) \quad j = 161, \dots, 240, \quad (5.9)$$

with all other elements of \mathbf{l}_k and \mathbf{f}_k set to zero for $k = 1, 2, 3$. This example has a sparse bi-cluster structure where distinct groups of samples are each loaded on only one factor (Figure 2a), and both the size of the groups and number of variables in each factor vary.

We applied *flash*, softImpute, SSVD and PMD to this example. (We excluded SFA and SFAmix since these methods do not model sparsity in both factors and loadings.) The



(a) **Left:** Illustration of the true latent rank-3 block structure used in these simulations. **Right** boxplots comparing accuracy of *flash* with several other methods across 100 replicates. Accuracy is measured by the difference of each methods RRMSE from the *flash* RRMSE, so smaller is better.



(b) Illustration of tendency of each method to either over-shrink the signal (SSVD) or under-shrink the noise (SI.cv, PMD.cv1, SVD) compared with *flash*. Each panel shows the mean absolute value of the estimated structure from each method.

Figure 2: Results from simulations with sparse bi-cluster structure ($K = 3$).

results (Figure 2) show that again *flash* consistently outperforms the other methods, and again the next best is softImpute. On this example both SSVD and PMD outperform SVD. Although SSVD and PMD perform similarly on average, their qualitative behavior is different: PMD insufficiently shrink the 0 values, whereas SSVD shrinks the 0 values well but overshrinks some of the signal, essentially removing the smallest of the three loading/factor combinations (Figure 2b).

5.2 Missing Data Imputation for Real Data Sets

Here we compare methods in their ability to impute missing data using five real data sets. In each case we “hold out” (mask) some of the data points, and then apply the methods to obtain estimates of the missing values. The data sets are as follows:

MovieLens 100K data, an (incomplete) 943×1682 matrix of user-movie ratings (integers from 1 to 5) (Harper and Konstan, 2016). Most users do not rate most movies, so the matrix is sparsely observed (94% missing), and contains about 100K observed ratings. We hold out a fraction of the observed entries and assess accuracy of methods in estimating these. We centered and scaled the ratings for each user before analysis.

GTEX eQTL summary data, a $16\,069 \times 44$ matrix of Z scores computed testing association of genetic variants (rows) with gene expression in different human tissues (columns). These data come from the Genotype Tissue Expression (GTEx) project (Consortium et al., 2015), which assessed the effects of thousands of “eQTLs” across 44 human tissues. (An eQTL is a genetic variant that is associated with expression of a gene.) To identify eQTLs, the project tested for association between expression and every near-by genetic variant, each test yielding a Z score. The data used here are the Z scores for the most significant genetic variant for each gene (the “top” eQTL). See Section 5.3 for more detailed analyses of these data.

Brain Tumor data, a 43×356 matrix of gene expression measurements on 4 different types of brain tumor (included in the `denoiseR` package, Josse et al., 2018). We centered each column before analysis.

Presidential address data, a 13×836 matrix of word counts from the inaugural addresses of 13 US presidents (1940–2009) (also included in the `denoiseR` package, Josse et al., 2018). Since both row and column means vary greatly we centered and scaled both rows and columns before analysis, using the `biScale` function from `softImpute`.

Breast cancer data, a 251×226 matrix of gene expression measurements from Carvalho et al. (2008), which were used as an example in the paper introducing NBSFA (Knowles and Ghahramani, 2011). Following Knowles and Ghahramani (2011) we centered each column (gene) before analysis.

Among the methods considered above, only *flash*, PMD and `softImpute` can handle missing data. We add NBSFA (Knowles and Ghahramani, 2011) to these comparisons. To emphasize the importance of parameter tuning we include results for PMD and `softImpute` with default settings (denoted PMD, SI) as well as using cross-validation (PMD.cv1, SI.cv).

For these real data the appropriate value of K is, of course, unknown. Both *flash* and NBSFA automatically estimate K . For PMD and `softImpute` we specified K based on the values inferred by *flash* and NBSFA. (Specifically, we used $K = 10, 30, 20, 10, 40$ respectively for the five data sets.)

We applied each method to all 5 data sets, using 10-fold OCV (Appendix B) to mask data points for imputation, repeated 20 times (with different random number seeds) for each data set. We measure imputation accuracy using root mean squared error (RMSE):

$$\text{RMSE}(\hat{Y}, Y; \Omega) = \sqrt{\frac{1}{|\Omega|} \sum_{ij \in \Omega} (Y_{ij} - \hat{Y}_{ij})^2}. \quad (5.10)$$

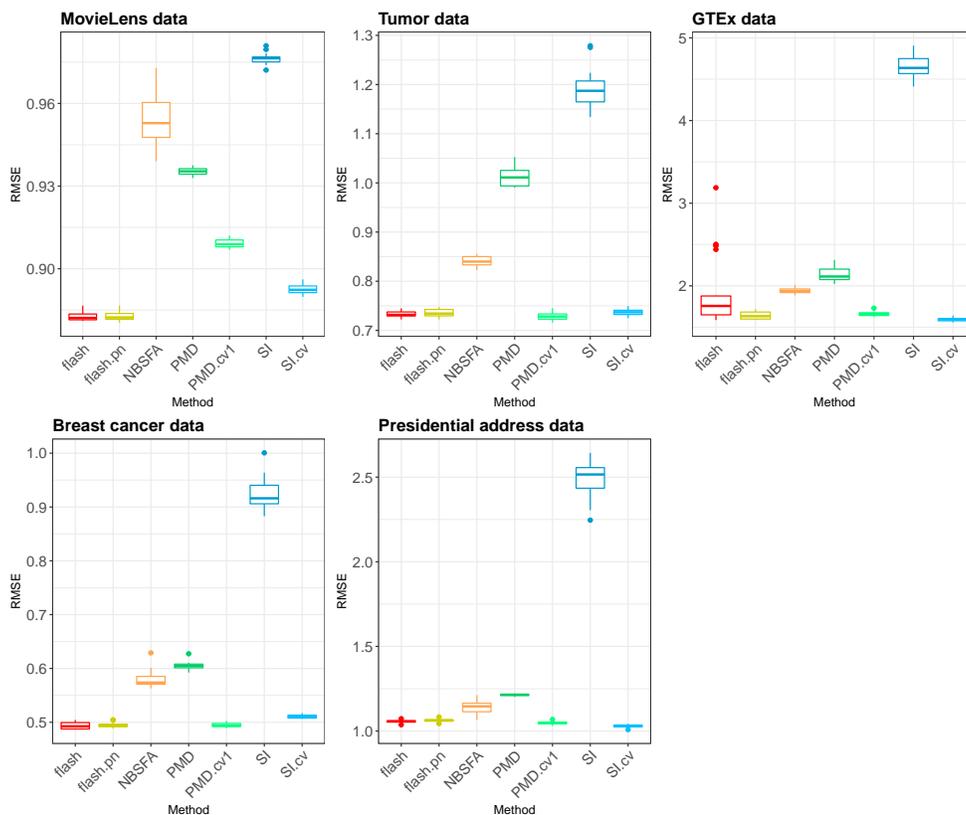


Figure 3: Comparison of the accuracy of different methods in imputing missing data. Each panel shows a boxplot of error rates (RMSE) for 20 simulations based on masking observed entries in a real data set.

where Ω is the set of indices of the held-out data points.

The results are shown in Figure 3. Although the ranking of methods varies among data sets, *flash*, PMD.cv1 and SI.cv perform similarly on average, and consistently outperform NBSFA, which in turn typically outperforms (untuned) PMD and unpenalized softImpute. These results highlight the importance of appropriate tuning for the penalized methods, and also the effectiveness of the EB method in *flash* to provide automatic tuning.

In these comparisons, as in the simulations, the two *flash* methods typically performed similarly. The exception is the GTEx data, where the scale mixture of normals ($\mathcal{G} = SN$) performed worse. Detailed investigation revealed this to be due to a very small number of very large “outlier” imputed values, well outside the range of the observed data, which grossly inflated RMSE. These outliers were so extreme that it should be possible to implement a filter to avoid them. However, we did not do this here as it seems useful to highlight this unexpected behavior. (Note that this occurs only when data are missing, and even then only in one of the five data sets considered here.)

5.3 Sharing of Genetic Effects on Gene Expression Among Tissues

To illustrate *flash* in a scientific application, we applied it to the GTEx data described above, a $16,069 \times 44$ matrix of Z scores, with Z_{ij} reflecting the strength (and direction) of effect of eQTL i in tissue j . We applied *flash* with $\mathcal{G} = \mathcal{SN}$ using the greedy+backfitting algorithm (i.e. the backfitting algorithm, initialized using the greedy algorithm).

The *flash* results yielded 26 factors (Figure 4-5) which summarize the main patterns of eQTL sharing among tissues (and, conversely, the main patterns of tissue-specificity). For example, the first factor has approximately equal weight for every tissue, and reflects the fact that many eQTLs show similar effects across all 44 tissues. The second factor has strong effects only in the 10 brain tissues, from which we infer that some eQTLs show much stronger effects in brain tissues than other tissues.

Subsequent factors tend to be sparser, and many have a strong effect in only one tissue, capturing “tissue-specific” effects. For example, the 3rd factor shows a strong effect only in whole blood, and captures eQTLs that have much stronger effects in whole blood than other tissues. (Two tissues, “Lung” and “Spleen”, show very small effects in this factor but with the same sign as blood. This is intriguing since the lung has recently been found to make blood cells—see Lefrançois et al. 2017—and a key role of the spleen is storing of blood cells.) Similarly Factors 7, 11 and 14 capture effects specific to “Testis”, “Thyroid” and “Esophagus Mucosa” respectively.

A few other factors show strong effects in a small number of tissues that are known to be biologically related, providing support that the factors identified are scientifically meaningful. For example, factor 10 captures the two tissues related to the cerebellum, “Brain Cerebellar Hemisphere” and “Brain Cerebellum”. Factor 19 captures tissues related to female reproduction, “Ovary”, “Uterus” and “Vagina”. Factor 5 captures “Muscle Skeletal”, with small but concordant effects in the heart tissues (“Heart Atrial Appendage” and “Heart Left Ventricle”). Factor 4, captures the two skin tissues (“Skin Not Sun Exposed Suprapubic”, “Skin Sun Exposed Lower leg”) and also “Esophagus Mucosa”, possibly reflecting the sharing of squamous cells that are found in both the surface of the skin, and the lining of the digestive tract. In factor 24, “Colon Transverse” and “Small Intestine Terminal Ileum” show the strongest effects (and with same sign), reflecting some sharing of effects in these intestinal tissues. Among the 26 factors, only a few are difficult to interpret biologically (e.g. factor 8).

To highlight the benefits of sparsity, we contrast the *flash* results with those for softImpute, which was the best-performing method in the missing data assessments on these data, but which uses a nuclear norm penalty that does not explicitly reward sparse factors or loadings. The first eight softImpute factors are shown in Figure 6. The softImpute results—except for the first two factors—show little resemblance to the flash results, and in our view are harder to interpret.

5.4 Computational Demands

It is difficult to make general statements about computational demands of our methods, because both the number of factors and number of iterations per factor can vary considerably depending on the data. However, to give a specific example, running our current implementation of the greedy algorithm on the GTEx data (a 16,000 by 44 matrix) takes

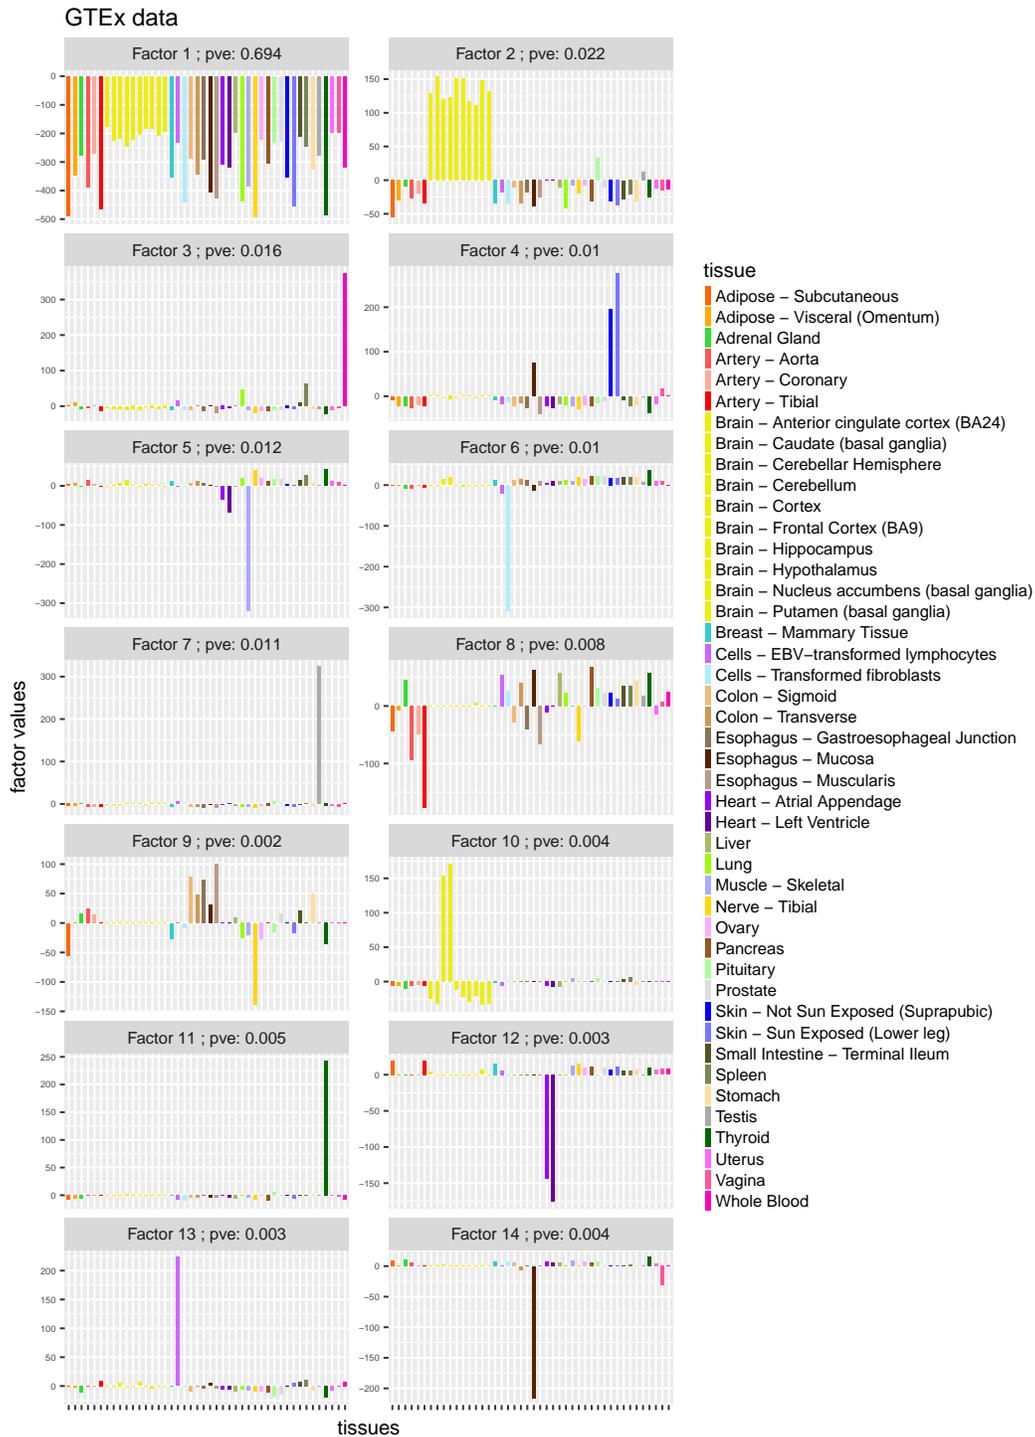


Figure 4: Results from running *flash* on GTEx data (factors 1 - 8). The pve (“Percentage Variance Explained”) for loading/factor k is defined as $pve_k := s_k / (\sum_k s_k + \sum_{ij} 1/\tau_{ij})$ where $s_k := \sum_{ij} (\bar{l}_{ki} f_{kj})^2$. It is a measure of the amount of signal in the data captured by loading/factor k (but its naming as “percentage variance explained” should be considered loose since the factors are not orthogonal).

EMPIRICAL BAYES MATRIX FACTORIZATION

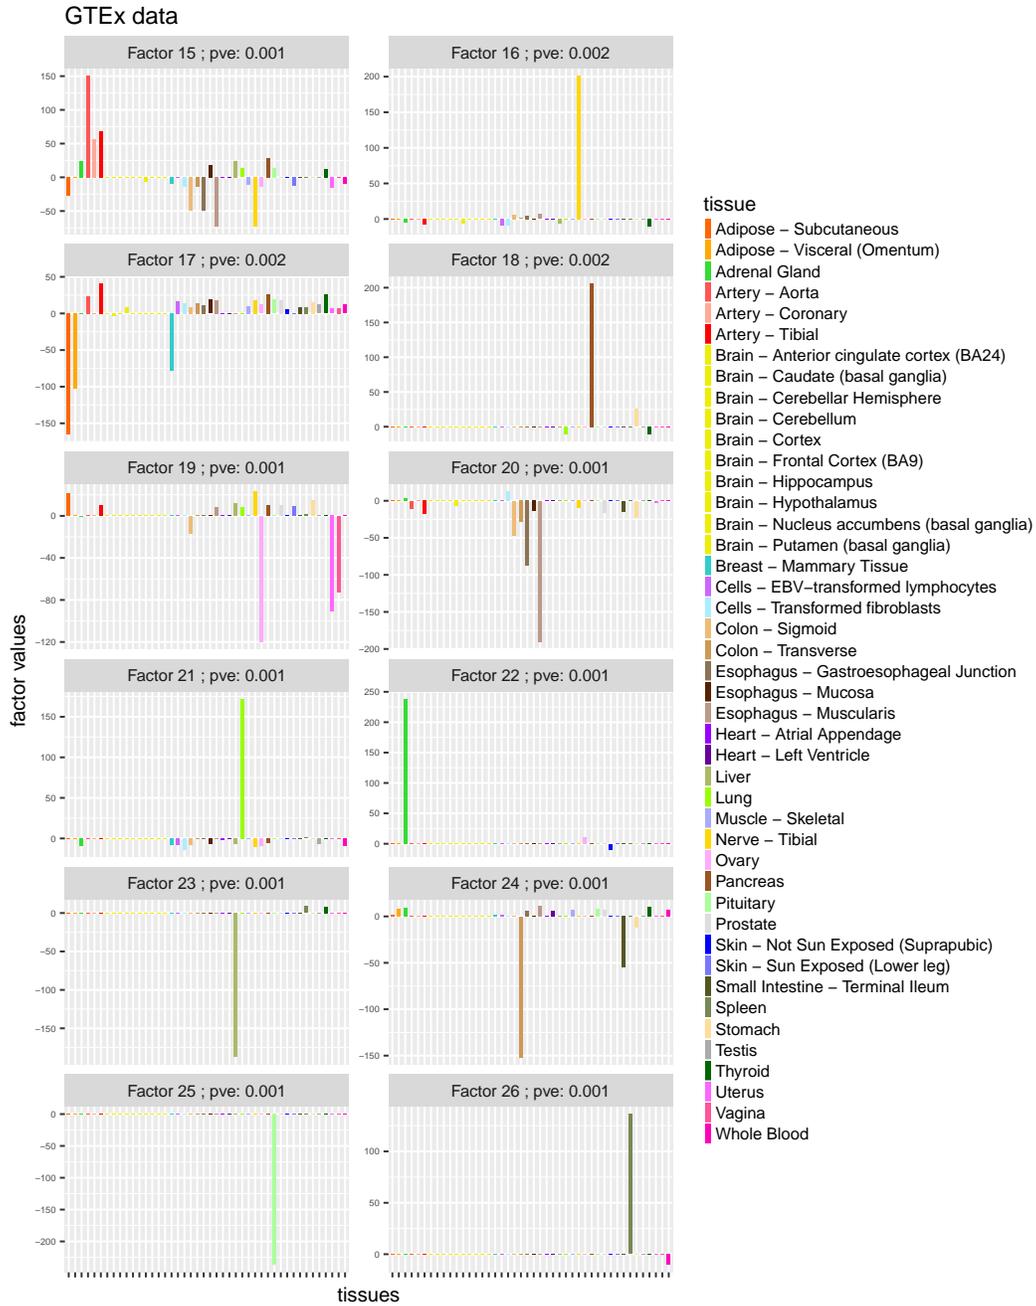


Figure 5: Results from running *flash* on GTEx data (factors 15 - 26)

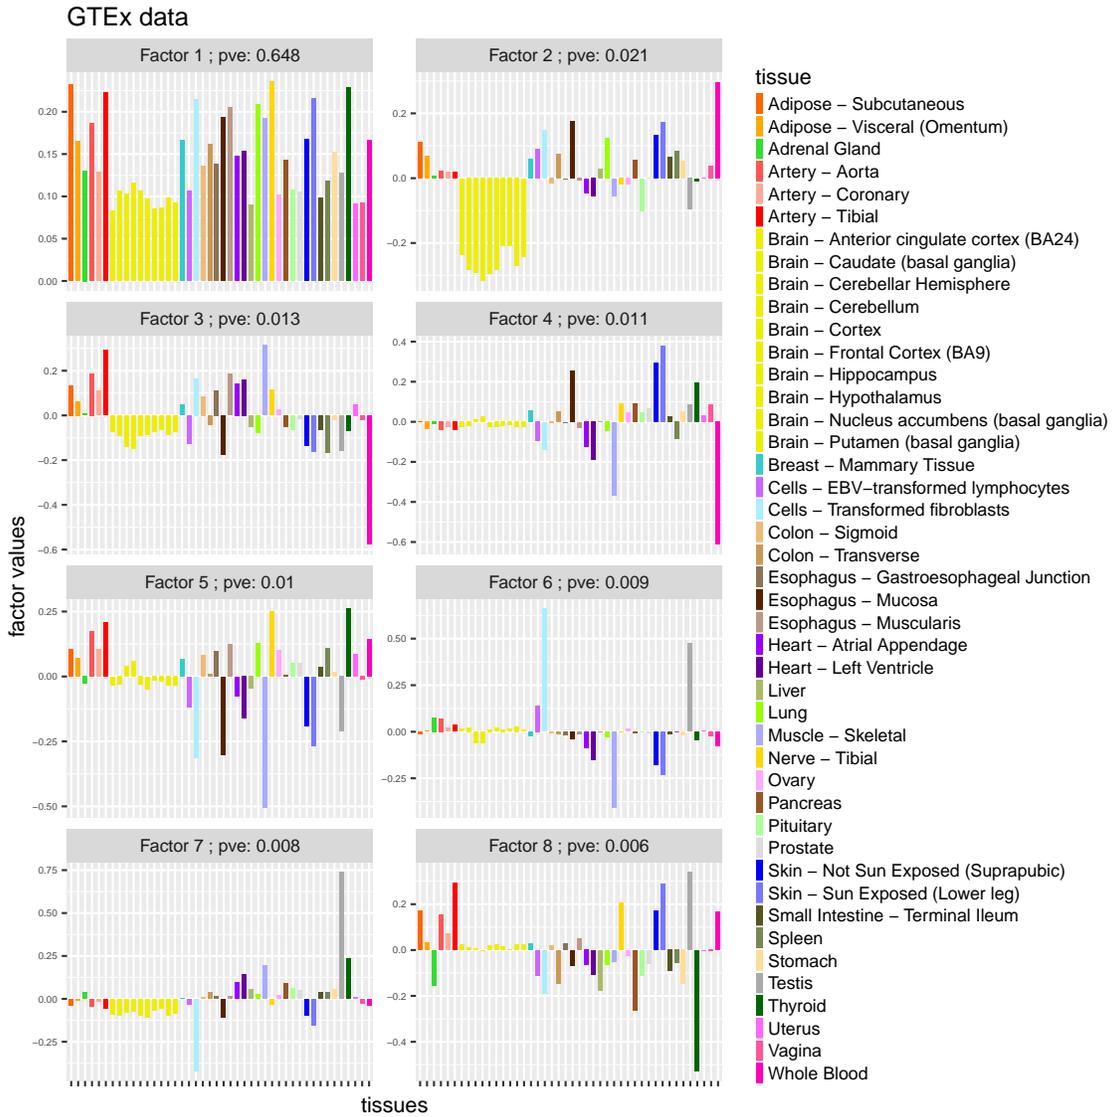


Figure 6: Results from running softImpute on GTEx data (factors 1-8). The factors are both less sparse and less interpretable than the *flash* results.

about 140s (wall time) for $\mathcal{G} = \mathcal{PN}$ and 650s for $\mathcal{G} = \mathcal{SN}$ (on a 2015 MACBOOK AIR with a 2.2 GHz Intel Core i7 processor and 8Gb RAM). By comparison, a single run of softImpute without CV takes 2-3s, so a naive implementation of 5-fold CV with 10 different tuning parameters and 10 different values of K would take over 1000s (although one could improve on this by use of warm starts for example).

6. Discussion

Here we discuss some potential extensions or modifications of our work.

6.1 Orthogonality Constraint

Our formulation here does not require the factors or loadings to be orthogonal. In scientific applications we do not see any particular reason to expect underlying factors to be orthogonal. However, imposing such a constraint could have computational or mathematical advantages. Formally adding such a constraint to our objective function seems tricky, but it would be straightforward to modify our algorithms to include an orthogonalization step each update. This would effectively result in an EB version of the SSVD algorithms in Yang et al. (2014), and it seems likely to be computationally faster than our current approach. One disadvantage of this approach is that it is unclear what optimization problem such an algorithm would solve (but the same is true of SSVD, and our algorithms have the advantage that they deal with missing data.)

6.2 Non-negative Matrix Factorization

We focused here on the potential for EBMF to induce sparsity on loadings and factors. However, EBMF can also encode other assumptions. For example, to assume the loadings and factors are non-negative, simply restrict \mathcal{G} to be a family of non-negative-valued distributions, yielding “Empirical Bayes non-negative Matrix Factorization” (EBNMF). Indeed, the `ashr` software can already solve the EBNM problem for some such families \mathcal{G} , and so `flash` already implements EBNMF. In preliminary assessments we found that the greedy approach is problematic here: the non-negative constraint makes it harder for later factors to compensate for errors in earlier factors. However, it is straightforward to apply the backfitting algorithm to fit EBNMF, with initialization by any existing NMF method. The performance of this approach is an area for future investigation.

6.3 Tensor Factorization

It is also straightforward to extend EBMF to tensor factorization, specifically a CANDECOMP/PARAFAC decomposition (Kolda and Bader, 2009):

$$Y_{ijm} = \sum_{k=1}^K l_{ki} f_{kj} h_{km} + E_{ijm} \quad (6.1)$$

$$l_{k1}, \dots, l_{kn} \sim^{iid} g_{\mathbf{l}_k}, \quad g_{\mathbf{l}_k} \in \mathcal{G} \quad (6.2)$$

$$f_{k1}, \dots, f_{kp} \sim^{iid} g_{\mathbf{f}_k}, \quad g_{\mathbf{f}_k} \in \mathcal{G} \quad (6.3)$$

$$h_{k1}, \dots, h_{kr} \sim^{iid} g_{\mathbf{h}_k}, \quad g_{\mathbf{h}_k} \in \mathcal{G} \quad (6.4)$$

$$E_{ijm} \sim^{iid} N(0, 1/\tau_{ijm}). \quad (6.5)$$

The variational approach is easily extended to this case (a generalization of methods in Hore et al., 2016), and updates that increase the objective function can be constructed by solving an EBNM problem, similar to EBMF. It seems likely that issues of convergence to local optima, and the need for good initializations, will need some attention to obtain good practical performance. However, results in Hore et al. (2016) are promising, and the automatic-tuning feature of EB methods seems particularly attractive here. For example, extending PMD to this case—allowing for different sparsity levels in l , f and h —would require 3 penalty parameters even in the rank 1 case, making it difficult to tune by CV.

6.4 Non-Gaussian Errors

It is also possible to extend the variational approximations used here to fit non-Gaussian models, such as binomial data; see for example Jaakkola and Jordan (2000); Seeger and Bouchard (2012); Klami (2015). The extension of our EB methods using these ideas is detailed in Wang (2017).

Acknowledgments

We thank P. Carbonetto for computational assistance, and P. Carbonetto, D. Gerard, and A. Sarkar for helpful conversations and comments on a draft manuscript. Computing resources were provided by the University of Chicago Research Computing Center. This work was supported by NIH grant HG002585 and by a grant from the Gordon and Betty Moore Foundation (Grant GBMF #4559).

Appendix A. Variational EBMF with K Factors

Here we describe in detail the variational approach to the K factor model, including deriving updates that we use to optimize the variational objective. (These derivations naturally include the $K = 1$ model as a special case, and our proof of Proposition 5 below includes Proposition 2 as a special case.)

Let $q_{\mathbf{l}}, q_{\mathbf{f}}$ denote the variational distributions on the K loadings/factors:

$$q_{\mathbf{l}}(\mathbf{l}_1, \dots, \mathbf{l}_K) = \prod_k q_{\mathbf{l}_k}(\mathbf{l}_k) \quad (\text{A.1})$$

$$q_{\mathbf{f}}(\mathbf{f}_1, \dots, \mathbf{f}_K) = \prod_k q_{\mathbf{f}_k}(\mathbf{f}_k). \quad (\text{A.2})$$

The objective function F (3.4) is thus a function of $q_{\mathbf{l}} = (q_{\mathbf{l}_1}, \dots, q_{\mathbf{l}_K})$, $q_{\mathbf{f}} = (q_{\mathbf{f}_1}, \dots, q_{\mathbf{f}_K})$, $g_{\mathbf{l}} = (g_{\mathbf{l}_1}, \dots, g_{\mathbf{l}_K})$ and $g_{\mathbf{f}} = (g_{\mathbf{f}_1}, \dots, g_{\mathbf{f}_K})$, as well as the precision $\boldsymbol{\tau}$:

$$F(q_{\mathbf{l}}, q_{\mathbf{f}}, g_{\mathbf{l}}, g_{\mathbf{f}}, \boldsymbol{\tau}) = \int \prod_k q_{\mathbf{l}_k}(\mathbf{l}_k) q_{\mathbf{f}_k}(\mathbf{f}_k) \log \frac{p(Y, \mathbf{l}, \mathbf{f}; g_{\mathbf{l}_1}, g_{\mathbf{f}_1}, \dots, g_{\mathbf{l}_K}, g_{\mathbf{f}_K}, \boldsymbol{\tau})}{\prod_k q_{\mathbf{l}_k}(\mathbf{l}_k) q_{\mathbf{f}_k}(\mathbf{f}_k)} d\mathbf{l}_k d\mathbf{f}_k, \quad (\text{A.3})$$

$$= E_{q_{\mathbf{l}}, q_{\mathbf{f}}} \log p(Y | \mathbf{l}, \mathbf{f}; \boldsymbol{\tau}) + \sum_k E_{q_{\mathbf{l}_k}} \log \frac{g_{\mathbf{l}_k}(\mathbf{l}_k)}{q_{\mathbf{l}_k}(\mathbf{l}_k)} + \sum_k E_{q_{\mathbf{f}_k}} \log \frac{g_{\mathbf{f}_k}(\mathbf{f}_k)}{q_{\mathbf{f}_k}(\mathbf{f}_k)}. \quad (\text{A.4})$$

We optimize F by iteratively updating parameters relating to $\boldsymbol{\tau}$, a single loading k ($q_{\mathbf{l}_k}, g_{\mathbf{l}_k}$) or factor k ($q_{\mathbf{f}_k}, g_{\mathbf{f}_k}$), keeping other parameters fixed. We simplify implementation by keeping track of only the first and second moments of the distributions $q_{\mathbf{l}_k}$ and $q_{\mathbf{f}_k}$, which we denote $\bar{\mathbf{l}}_k, \bar{\mathbf{l}}_k^2, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_k^2$. We now describe each kind of update in turn.

A.1 Updates for Precision Parameters

Here we derive updates to optimize F over the precision parameters $\boldsymbol{\tau}$. Focusing on the parts of F that depend on $\boldsymbol{\tau}$ gives:

$$F(\boldsymbol{\tau}) = E_{q_{\mathbf{l}}} E_{q_{\mathbf{f}}} \sum_{ij} 0.5 \log(\tau_{ij}) - 0.5 \tau_{ij} (Y_{ij} - \sum_k l_{ki} f_{kj})^2 + \text{const} \quad (\text{A.5})$$

$$= 0.5 \sum_{ij} \left[\log(\tau_{ij}) + \tau_{ij} \bar{R}_{ij}^2 \right] + \text{const} \quad (\text{A.6})$$

where \bar{R}^2 is defined by:

$$\bar{R}_{ij}^2 := E_{q_{\mathbf{l}}, q_{\mathbf{f}}} [(Y_{ij} - \sum_{k=1}^K l_{ki} f_{kj})^2] \quad (\text{A.7})$$

$$= (Y_{ij} - \sum_k \bar{l}_{ki} \bar{f}_{kj})^2 - \sum_k (\bar{l}_{ki})^2 (\bar{f}_{kj})^2 + \sum_k \bar{l}_{ki}^2 \bar{f}_{kj}^2. \quad (\text{A.8})$$

If we constrain $\boldsymbol{\tau} \in \mathcal{T}$ then we have

$$\hat{\boldsymbol{\tau}} = \arg \max_{\boldsymbol{\tau} \in \mathcal{T}} \sum_{ij} [\log(\tau_{ij}) - \tau_{ij} \overline{R^2}_{ij}]. \quad (\text{A.9})$$

For example, assuming constant precision $\tau_{ij} = \boldsymbol{\tau}$ yields:

$$\hat{\boldsymbol{\tau}} = \frac{NP}{\sum_{ij} \overline{R^2}_{ij}}. \quad (\text{A.10})$$

Assuming column-specific precisions ($\tau_{ij} = \boldsymbol{\tau}_j$), which is the default in our software, yields:

$$\hat{\boldsymbol{\tau}}_j = \frac{N}{\sum_i \overline{R^2}_{ij}}. \quad (\text{A.11})$$

Other variance structures are considered in Appendix A.5 of Wang (2017).

A.2 Updating Loadings and Factors

The following Proposition, which generalizes Proposition 2 in the main text, shows how updates for loadings (and factors) for the K -factor EBMF model can be achieved by solving an EBNM problem.

Proposition 5 *For the K -factor model, $\arg \max_{q_l, q_f, g_l, g_f, \boldsymbol{\tau}} F(q_l, q_f, g_l, g_f, \boldsymbol{\tau})$ is solved by solving an EBNM problem. Specifically*

$$\arg \max_{q_l, q_f, g_l, g_f, \boldsymbol{\tau}} F(q_l, q_f, g_l, g_f, \boldsymbol{\tau}) = \text{EBNM}(\hat{\boldsymbol{l}}(R^k, \bar{\boldsymbol{f}}_k, \bar{\boldsymbol{f}}^2_k, \boldsymbol{\tau}), \boldsymbol{s}_l(\bar{\boldsymbol{f}}^2_k, \boldsymbol{\tau})) \quad (\text{A.12})$$

where the functions $\hat{\boldsymbol{l}}$ and \boldsymbol{s}_l are given by (3.19) and (3.20), $\bar{\boldsymbol{f}}_k, \bar{\boldsymbol{f}}^2_k \in R^p$ denote the vectors whose elements are the first and second moments of \boldsymbol{f}_k under $q_{\boldsymbol{f}_k}$, and R^k denotes the residual matrix (3.27).

Similarly, $\arg \max_{q_{\boldsymbol{f}_k}, g_{\boldsymbol{f}_k}} F(q_l, q_f, g_l, g_f, \boldsymbol{\tau})$ is solved by solving an EBNM problem. Specifically,

$$\arg \max_{q_{\boldsymbol{f}_k}, g_{\boldsymbol{f}_k}} F(q_l, q_f, g_l, g_f, \boldsymbol{\tau}) = \text{EBNM}(\hat{\boldsymbol{f}}(R^k, \bar{\boldsymbol{l}}_k, \bar{\boldsymbol{l}}^2_k, \boldsymbol{\tau}), \boldsymbol{s}_f(\bar{\boldsymbol{l}}^2_k, \boldsymbol{\tau})) \quad (\text{A.13})$$

where the functions $\hat{\boldsymbol{f}} : R^{n \times p} \times R^n \times R^n \times R^{n \times p} \rightarrow R^p$ and $\boldsymbol{s}_f : R^n \times R^{n \times p} \rightarrow R^p$ are given by

$$\hat{\boldsymbol{f}}(Y, \boldsymbol{v}, \boldsymbol{w}, \boldsymbol{\tau})_j := \frac{\sum_i \tau_{ij} Y_{ij} v_i}{\sum_i \tau_{ij} w_i}, \quad (\text{A.14})$$

$$\boldsymbol{s}_f(\boldsymbol{w}, \boldsymbol{\tau})_j := \left(\sum_i \tau_{ij} w_i \right)^{-0.5}. \quad (\text{A.15})$$

A.2.1 A LEMMA ON THE NORMAL MEANS PROBLEM

To prove Proposition 5 we introduce a lemma that characterizes the solution of the normal means problem in terms of an objective that is closely related to the variational objective.

Recall that the EBNM model is:

$$\mathbf{x} = \boldsymbol{\theta} + \mathbf{e} \quad (\text{A.16})$$

$$\theta_1, \dots, \theta_n \sim^{iid} g, \quad g \in \mathcal{G}. \quad (\text{A.17})$$

where $e_i \sim N(0, s_i^2)$.

Solving the EBNM problem involves estimating g by maximum likelihood:

$$\hat{g} = \arg \max_{g \in \mathcal{G}} l(g), \quad (\text{A.18})$$

where

$$l(g) = \log p(\mathbf{x}|g). \quad (\text{A.19})$$

It also involves finding the posterior distributions:

$$p(\boldsymbol{\theta}|\mathbf{x}, \hat{g}) = \prod_j p(\theta_j|\mathbf{x}, \hat{g}) \propto \prod_j \hat{g}(\theta_j) p(x_j|\theta_j, s_j). \quad (\text{A.20})$$

Lemma 6 *Solving the EBNM problem also solves:*

$$\max_{q_\theta, g \in \mathcal{G}} F^{NM}(q_\theta, g) \quad (\text{A.21})$$

where

$$F^{NM}(q_\theta, g) = E_{q_\theta} \left[-\frac{1}{2} \sum_j (A_j \theta_j^2 - 2B_j \theta_j) \right] + E_{q_\theta} \log \frac{g(\boldsymbol{\theta})}{q_\theta(\boldsymbol{\theta})} + \text{const} \quad (\text{A.22})$$

with $A_j = 1/s_j^2$ and $B_j = x_j/s_j^2$, and $g(\boldsymbol{\theta}) := \prod_j g(\theta_j)$.

Equivalently, (A.21)-(A.22) is solved by $g = \hat{g}$ in (A.18) and $q_\theta = p(\boldsymbol{\theta}|\mathbf{x}, \hat{g})$ in (A.20), with $x_j = B_j/A_j$ and $s_j^2 = 1/A_j$.

Proof The log likelihood can be written as

$$l(g) := \log[p(\mathbf{x}|g)] \quad (\text{A.23})$$

$$= \log[p(\mathbf{x}, \boldsymbol{\theta}|g)/p(\boldsymbol{\theta}|\mathbf{x}, g)] \quad (\text{A.24})$$

$$= \int q_\theta(\boldsymbol{\theta}) \log \frac{p(\mathbf{x}, \boldsymbol{\theta}|g)}{p(\boldsymbol{\theta}|\mathbf{x}, g)} d\boldsymbol{\theta} \quad (\text{A.25})$$

$$= \int q_\theta(\boldsymbol{\theta}) \log \frac{p(\mathbf{x}, \boldsymbol{\theta}|g)}{q_\theta(\boldsymbol{\theta})} d\boldsymbol{\theta} + \int q_\theta(\boldsymbol{\theta}) \log \frac{q_\theta(\boldsymbol{\theta})}{p(\boldsymbol{\theta}|\mathbf{x}, g)} d\boldsymbol{\theta} \quad (\text{A.26})$$

$$= F^{NM}(q_\theta, g) + D_{KL}(q_\theta || p_{\theta|\mathbf{x},g}) \quad (\text{A.27})$$

where

$$F^{NM}(q_\theta, g) = \int q_\theta(\boldsymbol{\theta}) \log \frac{p(\mathbf{x}, \boldsymbol{\theta}|g)}{q_\theta(\boldsymbol{\theta})} d\boldsymbol{\theta} \quad (\text{A.28})$$

and

$$D_{KL}(q_\theta || p_{\theta|x,g}) = - \int q_\theta(\boldsymbol{\theta}) \log \frac{p(\boldsymbol{\theta}|\mathbf{x},g)}{q_\theta(\boldsymbol{\theta})} d\boldsymbol{\theta} \quad (\text{A.29})$$

Here $p_{\theta|x,g}$ denotes the posterior distribution $p(\boldsymbol{\theta}|\mathbf{x},g)$. This identity holds for any distribution $q_\theta(\boldsymbol{\theta})$.

Rearranging (A.27) gives:

$$F^{\text{NM}}(q_\theta, g) = l(g) - D_{KL}(q_\theta || p_{\theta|x,g}). \quad (\text{A.30})$$

Since $D_{KL}(q_\theta || p_{\theta|x,g}) \geq 0$, with equality when $q_\theta = p_{\theta|x,g}$, $F^{\text{NM}}(q_\theta, g)$ is maximized over q_θ by setting $q_\theta = p_{\theta|x,g}$. Further

$$\max_{q_\theta} F^{\text{NM}}(q_\theta, g) = l(g), \quad (\text{A.31})$$

so

$$\arg \max_{g \in \mathcal{G}} \max_{q_\theta} F^{\text{NM}}(q_\theta, g) = \arg \max_{g \in \mathcal{G}} l(g) = \hat{g}. \quad (\text{A.32})$$

It remains only to show that F^{NM} has the form (A.22).

By (A.16) and (A.17), we have

$$\log p(\mathbf{x}, \boldsymbol{\theta}|g) = -\frac{1}{2} \sum_j s_j^{-2} (x_j - \theta_j)^2 + \log g(\boldsymbol{\theta}) + \text{const}. \quad (\text{A.33})$$

Thus

$$F^{\text{NM}}(q_\theta, g) = E_{q_\theta} \left[-\frac{1}{2} \sum_j (A_j \theta_j^2 - 2B_j \theta_j) \right] + E_{q_\theta} \log \frac{g(\boldsymbol{\theta})}{q_\theta(\boldsymbol{\theta})} + \text{const}. \quad (\text{A.34})$$

■

A.2.2 PROOF OF PROPOSITION 5

We are now ready to prove Proposition 5.

Proof We prove the first part of the proposition since the proof for the second part is essentially the same.

The objective function (A.3) is:

$$F(q_l, q_f, g_l, g_f, \boldsymbol{\tau}) = E_{q_l, q_f} \log p(Y|\mathbf{l}, \mathbf{f}; \boldsymbol{\tau}) + \sum_k E_{q_{l_k}} \log \frac{g_{l_k}(\mathbf{l}_k)}{q_{l_k}(\mathbf{l}_k)} + \sum_k E_{q_{f_k}} \log \frac{g_{f_k}(\mathbf{f}_k)}{q_{f_k}(\mathbf{f}_k)} \quad (\text{A.35})$$

$$= E_{q_{l_k}} \left[-\frac{1}{2} \sum_i (A_{ik} l_{ki}^2 - 2B_{ik} l_{ki}) \right] + E_{q_{l_k}} \log \frac{g_{l_k}(\mathbf{l}_k)}{q_{l_k}(\mathbf{l}_k)} + C_1 \quad (\text{A.36})$$

where C_1 is a constant with respect to q_k, g_k and

$$A_{ik} = \sum_j \tau_{ij} \mathbb{E}_{q_f}(f_{kj}^2) \quad (\text{A.37})$$

$$B_{ik} = \sum_j \tau_{ij} \left(R_{ij}^k \mathbb{E}_{q_f} f_{kj} \right). \quad (\text{A.38})$$

Based on Lemma 6, we can solve this optimization problem (A.36) by solving the EBNM problem with:

$$x_i = \frac{\sum_j \tau_{ij} \left(R_{ij}^k \mathbb{E}_{q_f} f_{kj} \right)}{\sum_j \tau_{ij} \mathbb{E}_{q_f}(f_{kj}^2)} \quad (\text{A.39})$$

$$s_i^2 = \frac{1}{\sum_j \tau_{ij} \mathbb{E}_{q_f}(f_{kj}^2)}. \quad (\text{A.40})$$

■

A.3 Algorithms

Just as with the rank 1 EBMF model, the updates for the rank K model require only the first and second moments of the variational distributions q . Thus we implement the updates in algorithms that keep track of the first moments ($\bar{\mathbf{l}} := (\bar{l}_1, \dots, \bar{l}_K)$ and $\bar{\mathbf{f}} := (\bar{\mathbf{f}}_1, \dots, \bar{\mathbf{f}}_K)$) and second moments ($\bar{\mathbf{l}}^2 := (\bar{l}_1^2, \dots, \bar{l}_K^2)$ and $\bar{\mathbf{f}}^2 := (\bar{\mathbf{f}}_1^2, \dots, \bar{\mathbf{f}}_K^2)$), and the precision $\boldsymbol{\tau}$.

Algorithm 3 implements a basic update for $\boldsymbol{\tau}$, and for the parameters relating to a single factor k ($\bar{\mathbf{l}}_k, \bar{l}_k^2, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_k^2$). Note that the latter updates are identical to the updates for fitting the single factor EBMF model, but with Y_{ij} replaced with the residuals obtained by removing the estimated effects of the other $k - 1$ factors.

Algorithm 3 Single-factor update for EBMF (rank K)

Require: A data matrix Y ($n \times p$)

Require: A function, $\text{ebnm}(\mathbf{x}, \mathbf{s}) \rightarrow (\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}}^2)$, that solves the EBNM problem (3.11)-(3.12) and returns the first and second posterior moments (3.15)-(3.16).

Require: Current values for first moments $\bar{\mathbf{l}} := (\bar{l}_1, \dots, \bar{l}_K)$ and $\bar{\mathbf{f}} := (\bar{f}_1, \dots, \bar{f}_K)$.

Require: Current values for second moments $\bar{\mathbf{l}}^2 := (\bar{l}_1^2, \dots, \bar{l}_K^2)$ and $\bar{\mathbf{f}}^2 := (\bar{f}_1^2, \dots, \bar{f}_K^2)$.

Require: An index k indicating which loading/factor to compute updated values for.

- 1: Compute matrix of expected squared residuals, \bar{R}^2 , using (A.7)
 - 2: $\tau_j \leftarrow n / \sum_i \bar{R}^2_{ij}$. [Assumes column-specific variances; can be modified to make other assumptions.]
 - 3: Compute residual matrix $R^k := Y - \sum_{k' \neq k} \bar{\mathbf{l}}_{k'} \bar{\mathbf{f}}_{k'}^T$.
 - 4: Compute $\hat{\mathbf{l}}(R^k, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_k^2, \boldsymbol{\tau})$ and its standard error $\mathbf{s}_l(\boldsymbol{\tau}, \bar{\mathbf{f}}_k^2)$, using (3.19) and (3.20).
 - 5: $(\bar{\mathbf{l}}_k, \bar{l}_k^2) \leftarrow \text{ebnm}(\hat{\mathbf{l}}, \mathbf{s}_l)$.
 - 6: Compute $\hat{\mathbf{f}}(R^k, \bar{\mathbf{l}}_k, \bar{l}_k^2, \boldsymbol{\tau})$ and its standard error $\mathbf{s}_f(\bar{\mathbf{l}}_k, \boldsymbol{\tau})$.
 - 7: $(\bar{\mathbf{f}}_k, \bar{f}_k^2) \leftarrow \text{ebnm}(\hat{\mathbf{f}}, \mathbf{s}_f)$.
 - 8: **return** updated values $\bar{\mathbf{l}}_k, \bar{l}_k^2, \bar{\mathbf{f}}_k, \bar{f}_k^2, \boldsymbol{\tau}$.
-

Based on these basic updates we implemented two algorithms for fitting the K -factor EBMF model: the greedy algorithm, and the backfitting algorithm, as follows.

A.3.1 GREEDY ALGORITHM

The greedy algorithm is a forward procedure that, at the k th step, adds new factors and loadings $\mathbf{l}_k, \mathbf{f}_k$ by optimizing over $q_{\mathbf{l}_k}, q_{\mathbf{f}_k}, g_{\mathbf{l}_k}, g_{\mathbf{f}_k}$ while keeping the distributions related to previous factors fixed. Essentially this involves fitting the single-factor model to the residuals obtained by removing previous factors. The procedure stops adding factors when the estimated new factors (or loadings) are identically zero. The algorithm as follows:

Algorithm 4 Greedy Algorithm for EBMF

Require: A data matrix Y ($n \times p$)

Require: A function, $\text{ebnm}(\mathbf{x}, \mathbf{s}) \rightarrow (\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}}^2)$, that solves the EBNM problem (3.11)-(3.12) and returns the first and second posterior moments (3.15)-(3.16).

Require: A function, $\text{init}(Y) \rightarrow (\bar{\mathbf{l}}; \bar{\mathbf{f}})$ that provides initial estimates for the loadings and factors (see Section 4.2).

Require: A function $\text{single.update}(Y, \bar{\mathbf{l}}, \bar{\mathbf{f}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}^2, k) \rightarrow (\bar{\mathbf{l}}_k, \bar{\mathbf{l}}_{k}^2, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_{k}^2, \boldsymbol{\tau})$ implementing Algorithm 3.

- 1: initialize $K \leftarrow 0$.
 - 2: **repeat**
 - 3: $K \leftarrow K + 1$.
 - 4: Compute residual matrix $R_{ij} = Y_{ij} - \sum_{k=1}^{K-1} \bar{\mathbf{l}}_{ki} \bar{\mathbf{f}}_{kj}$.
 - 5: Initialize first moments $(\bar{\mathbf{l}}_K, \bar{\mathbf{f}}_K) \leftarrow \text{init}(R)$.
 - 6: Initialize second moments by squaring first moments: $\bar{\mathbf{l}}_K^2 \leftarrow \bar{\mathbf{l}}_K^2; \bar{\mathbf{f}}_K^2 \leftarrow \bar{\mathbf{f}}_K^2$.
 - 7: **repeat**
 - 8: $(\bar{\mathbf{l}}_K, \bar{\mathbf{l}}_K^2, \bar{\mathbf{f}}_K, \bar{\mathbf{f}}_K^2, \boldsymbol{\tau}) \leftarrow \text{single.update}(Y, \bar{\mathbf{l}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, K)$
 - 9: **until** converged
 - 10: **until** $\bar{\mathbf{f}}_K$ is identically 0 or $\bar{\mathbf{l}}_K$ is identically 0.
 - 11: **return** $\bar{\mathbf{l}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, \boldsymbol{\tau}$
-

A.3.2 BACKFITTING ALGORITHM

The backfitting algorithm iteratively refines a fit of K factors and loadings, by updating them one at a time, at each update keeping the other loadings and factors fixed. The name comes from its connection with the backfitting algorithm in Breiman and Friedman (1985), specifically the fact that it involves iteratively re-fitting to residuals.

Algorithm 5 Backfitting algorithm for EBMF (rank K)

Require: A data matrix Y ($n \times p$)

Require: A function, $\text{ebnm}(\mathbf{x}, \mathbf{s}) \rightarrow (\bar{\boldsymbol{\theta}}, \bar{\boldsymbol{\theta}}^2)$, that solves the EBNM problem (3.11)-(3.12) and returns the first and second posterior moments (3.15)-(3.16).

Require: A function, $\text{init}(Y) \rightarrow (\mathbf{l}_1, \dots, \mathbf{l}_K; \mathbf{f}_1, \dots, \mathbf{f}_K)$ that provides initial estimates for the loadings and factors (e.g. the greedy algorithm from Appendix A.3.1, or a rank K SVD).

Require: A function $\text{single.update}(Y, \bar{\mathbf{l}}, \bar{\mathbf{f}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}^2, k) \rightarrow (\bar{\mathbf{l}}_k, \bar{\mathbf{l}}_{k}^2, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_{k}^2, \boldsymbol{\tau})$ implementing Algorithm 3.

- 1: Initialize first moments $(\bar{\mathbf{l}}_1, \dots, \bar{\mathbf{l}}_K; \bar{\mathbf{f}}_1, \dots, \bar{\mathbf{f}}_K) \leftarrow \text{init}(Y)$.
 - 2: Initialize second moments by squaring first moments: $\bar{\mathbf{l}}_k^2 \leftarrow \bar{\mathbf{l}}_k^2; \bar{\mathbf{f}}_k^2 \leftarrow \bar{\mathbf{f}}_k^2$. [alternatively the init function could provide these initial values].
 - 3: **repeat**
 - 4: **for** $k = 1, \dots, K$ **do**
 - 5: $(\bar{\mathbf{l}}_k, \bar{\mathbf{l}}_{k}^2, \bar{\mathbf{f}}_k, \bar{\mathbf{f}}_{k}^2, \boldsymbol{\tau}) \leftarrow \text{single.update}(Y, \bar{\mathbf{l}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, k)$
 - 6: **until** converged
 - 7: **return** $\bar{\mathbf{l}}, \bar{\mathbf{l}}^2, \bar{\mathbf{f}}, \bar{\mathbf{f}}^2, \boldsymbol{\tau}$
-

A.4 Objective Function Computation

The algorithms above all involve updates that will increase (or, at least, not decrease) the objective function $F(q_l, q_f, g_l, g_f, \boldsymbol{\tau})$. However, these updates do not require computing the objective function itself. In iterative algorithms it can be helpful to compute the objective function to monitor convergence (and as a check on implementation). In this subsection we describe how this can be done. In essence, this involves extending the solver of the EBNM problem to also return the value of the log-likelihood achieved in that problem (which is usually not difficult).

The objective function of the EBMF model is:

$$F(q_l, q_f, g_l, g_f, \boldsymbol{\tau}) = E_{q_l, q_f} \log p(Y|\mathbf{l}, \mathbf{f}; \boldsymbol{\tau}) + E_{q_l} \log \frac{g_l(\mathbf{l})}{q_l(\mathbf{l})} + E_{q_f} \log \frac{g_f(\mathbf{f})}{q_f(\mathbf{f})} \quad (\text{A.41})$$

The calculation of $E_{q_l, q_f} \log p(Y|\mathbf{l}, \mathbf{f}; \boldsymbol{\tau})$ is straightforward and $E_{q_l} \log \frac{g_l(\mathbf{l})}{q_l(\mathbf{l})}$ and $E_{q_f} \log \frac{g_f(\mathbf{f})}{q_f(\mathbf{f})}$ can be calculated using the log-likelihood of the EBNM model using the following Lemma 7.

Lemma 7 *Suppose \hat{g}, q solves the EBNM problem with data (\mathbf{x}, \mathbf{s}) :*

$$(\hat{g}, q) = \text{EBNM}(\mathbf{x}, \mathbf{s}), \quad (\text{A.42})$$

where $q := (q_1, \dots, q_n)$ are the estimated posterior distributions of the normal means parameters $\theta_1, \dots, \theta_n$. Then

$$E_q(\log(\prod_j \hat{g}(\theta_j) / \prod_j q_j(\theta_j))) = l(\hat{g}; \mathbf{x}, \mathbf{s}) + \frac{1}{2} \sum_j \log(2\pi s_j^2) + (1/s_j^2)(x_j^2 + E_q(\theta_j^2)) - 2x_j E_q(\theta_j) \quad (\text{A.43})$$

where $l(\hat{g}; \mathbf{x}, \mathbf{s})$ is the log of the likelihood for the normal means problem (A.27).

Proof We have from (A.28)

$$F^{\text{NM}}(q_\theta, \hat{g}) = \int q_\theta(\boldsymbol{\theta}) \log \frac{p(\mathbf{x}, \boldsymbol{\theta}|\hat{g})}{q_\theta(\boldsymbol{\theta})} d\boldsymbol{\theta} \quad (\text{A.44})$$

$$= \int q_\theta(\boldsymbol{\theta}) \log \frac{p(\mathbf{x}|\boldsymbol{\theta})\hat{g}(\boldsymbol{\theta})}{q_\theta(\boldsymbol{\theta})} d\boldsymbol{\theta} \quad (\text{A.45})$$

$$= E_q(\log(\prod_j \hat{g}(\theta_j) / \prod_j q_j(\theta_j))) - \frac{1}{2} E_q[\sum_j \log(2\pi s_j^2) + (1/s_j^2)(x_j - \theta_j)^2] \quad (\text{A.46})$$

And the result follows from noting that $F^{\text{NM}}(\hat{q}, \hat{g}) = l(\hat{g})$. ■

A.5 Inference with Penalty Term

Conceivably, in some settings one might like to encourage solutions to the EBMF problem be sparser than the maximum-likelihood estimates for g_l, g_f would produce. This could

be done by extending the EBMF model to introduce a penalty term on the distributions $g_{\mathbf{l}}, g_{\mathbf{f}}$ so that the maximum likelihood estimates are replaced by maximizing a penalized likelihood. We are not advocating for this approach, but it is straightforward given existing machinery, and so we document it here for completeness.

Let $h_{\mathbf{l}}(g_{\mathbf{l}})$ and $h_{\mathbf{f}}(g_{\mathbf{f}})$ denote penalty terms on $g_{\mathbf{l}}$ and $g_{\mathbf{f}}$, so the penalized log-likelihood would be:

$$\begin{aligned} l(g_{\mathbf{l}}, g_{\mathbf{f}}, \boldsymbol{\tau}) &:= \log[p(Y|g_{\mathbf{l}}, g_{\mathbf{f}}, \boldsymbol{\tau}^2)] + h_{\mathbf{l}}(g_{\mathbf{l}}) + h_{\mathbf{f}}(g_{\mathbf{f}}) \\ &= F(q, g_{\mathbf{l}}, g_{\mathbf{f}}, \boldsymbol{\tau}^2) + h_{\mathbf{l}}(g_{\mathbf{l}}) + h_{\mathbf{f}}(g_{\mathbf{f}}) + D_{KL}(q||p) \end{aligned} \quad (\text{A.47})$$

where $F(q, g_{\mathbf{l}}, g_{\mathbf{f}}, \boldsymbol{\tau}^2)$ and $D_{KL}(q||p)$ are defined in (3.4) and (3.5). And the corresponding penalized variational objective is:

$$\max F(q, g_{\mathbf{l}}, g_{\mathbf{f}}, \boldsymbol{\tau}^2) + h_{\mathbf{l}}(g_{\mathbf{l}}) + h_{\mathbf{f}}(g_{\mathbf{f}}). \quad (\text{A.48})$$

It is straightforward to modify the algorithms above to maximize this penalized objective: simply modify the EBNM solvers to solve a corresponding penalized normal means problem. That is, instead of estimating the prior g by maximum likelihood, the EBNM solver must now maximize the penalized log-likelihood:

$$\hat{g} = \arg \max_{g \in \mathcal{G}} l_{\text{EBNM}}(g) + h(g), \quad (\text{A.49})$$

where l_{EBNM} denote the log-likelihood for the EBNM problem. (The computation of the posterior distributions given \hat{g} is unchanged).

For example, the `ashr` software (Stephens, 2017) provides the option to include a penalty on g to encourage overestimation of the size of the point mass on zero. This penalty was introduced to ensure conservative behavior in False Discovery Rate applications of the normal means problem. It is unclear that such a penalty is desirable in the matrix factorization application. However, the above discussion shows that using this penalty (e.g. within the `ebnm` function used by the greedy or backfitting algorithms) can be thought of as solving a penalized version of the EBMF problem.

Appendix B. Orthogonal Cross Validation

Cross-validation assessments involving “holding out” (hiding) data from methods. Here we introduce a novel approach to selecting the data to be held out, which we call Orthogonal Cross Validation (OCV). Although not the main focus of our paper, we believe that OCV is a novel and appealing approach to selecting hold-out data for factor models, e.g. when using CV to select an appropriate dimension K for dimension reduction methods, as in Owen and Wang (2016).

Generic k -fold CV involves randomly dividing the data matrix into k parts and then, for each part, training methods on the other $k - 1$ parts before assessing error on that part, as in Algorithm 6.

Algorithm 6 k-fold CV

- 1: **procedure** K-FOLD CROSS VALIDATION
 - 2: randomly divide data matrix Y into $Y_{(1)}, \dots, Y_{(k)}$ with “hold-out” index $\Omega_{(1)}, \dots, \Omega_{(k)}$
 - 3: **for** $i = 1, \dots, k$ **do**
 - 4: take $Y_{(i)}$ as missing and run *flash*
 - 5: $\hat{Y}_{(i)} = E[Y_{\Omega_{(i)}} | Y_{-\Omega_{(i)}}]$
 - 6: $s_i^2 = \|\hat{Y}_{(i)} - Y_{\Omega_{(i)}}\|_2^2$
- return** RMSE: $score = \sqrt{\frac{\sum_k s_k^2}{NP}}$
-

The novel part of OCV is in how to choose the “hold-out” pattern. We randomly divide the columns and rows into k sets. and put these sets into k orthogonal parts, and then take all Y_{ij} with the chosen column and row indices as “hold-out” $Y_{(i)}$.

To illustrate this scheme, we take 3-fold CV as an example. We randomly divide the columns into 3 sets and the rows into 3 sets as well. The data matrix Y is divided into 9 partition (by row and column permutation):

$$Y = \begin{pmatrix} Y_{11} & Y_{12} & Y_{13} \\ Y_{21} & Y_{22} & Y_{23} \\ Y_{31} & Y_{32} & Y_{33} \end{pmatrix}$$

Then $Y_{(1)} = \{Y_{11}, Y_{22}, Y_{33}\}$, $Y_{(2)} = \{Y_{12}, Y_{23}, Y_{31}\}$ and $Y_{(3)} = \{Y_{13}, Y_{21}, Y_{32}\}$ are orthogonal to each other. Then the data matrix Y is marked as:

$$Y = \begin{pmatrix} Y_{(1)} & Y_{(2)} & Y_{(3)} \\ Y_{(3)} & Y_{(1)} & Y_{(2)} \\ Y_{(2)} & Y_{(3)} & Y_{(1)} \end{pmatrix}$$

In OCV, each fold k , $Y_{(k)}$ contains equally balanced part of data matrix and includes all the row and column indices. This ensures that all i 's and j 's are included into each $Y_{-(k)}$. In 3-fold OCV, we have:

$$Y = \begin{pmatrix} Y_{11} & Y_{12} & Y_{13} \\ Y_{21} & Y_{22} & Y_{23} \\ Y_{31} & Y_{32} & Y_{33} \end{pmatrix} = \begin{bmatrix} L^{(1)} \\ L^{(2)} \\ L^{(3)} \end{bmatrix} \times [F^{(1)} \quad F^{(2)} \quad F^{(3)}] + E \quad (\text{B.1})$$

$$= \begin{pmatrix} Y_{(1)} & Y_{(2)} & Y_{(3)} \\ Y_{(3)} & Y_{(1)} & Y_{(2)} \\ Y_{(2)} & Y_{(3)} & Y_{(1)} \end{pmatrix} = \begin{bmatrix} L^{(1)}F^{(1)} & L^{(1)}F^{(2)} & L^{(1)}F^{(3)} \\ L^{(2)}F^{(1)} & L^{(2)}F^{(2)} & L^{(2)}F^{(3)} \\ L^{(3)}F^{(1)} & L^{(3)}F^{(2)} & L^{(3)}F^{(3)} \end{bmatrix} + E \quad (\text{B.2})$$

where $Y_{(1)} = \{Y_{11}, Y_{22}, Y_{33}\}$, $Y_{(2)} = \{Y_{12}, Y_{23}, Y_{31}\}$ and $Y_{(3)} = \{Y_{13}, Y_{21}, Y_{32}\}$. We can see for each “hold-out” part, $Y_{(k)}$, $L^{(1)}, L^{(2)}, L^{(3)}$ and $F^{(1)}, F^{(2)}, F^{(3)}$ show up once and only once. In this sense the hold-out pattern is “balanced”.

References

- Ricard Argelaguet, Britta Velten, Damien Arno, Sascha Dietrich, Thorsten Zenz, John C Marioni, Florian Buettner, Wolfgang Huber, and Oliver Stegle. Multi-omics factor analysis—a framework for unsupervised integration of multi-omics data sets. *Molecular systems biology*, 14(6):e8124, 2018.
- Hagai Attias. Independent factor analysis. *Neural Computation*, 11(4):803–851, 1999.
- Jushan Bai and Serena Ng. *Large Dimensional Factor Analysis*. Now Publishers Inc, 2008.
- Anirban Bhattacharya and David B Dunson. Sparse Bayesian infinite factor models. *Biometrika*, pages 291–306, 2011.
- Christopher M. Bishop. Variational principal components. In *Ninth International Conference on Artificial Neural Networks (Conf. Publ. No. 470)*, volume 1, pages 509–514. IET, 1999.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.
- Guillaume Bouchard, Jason Naradowsky, Sebastian Riedel, Tim Rocktäschel, and Andreas Vlachos. Matrix and tensor factorization methods for natural language processing. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing: Tutorial Abstracts*, pages 16–18, Beijing, China, July 2015. Association for Computational Linguistics. doi: 10.3115/v1/P15-5005. URL <https://www.aclweb.org/anthology/P15-5005>.
- Leo Breiman and Jerome H. Friedman. Estimating optimal transformations for multiple regression and correlation. *Journal of the American Statistical Association*, 80(391):580–598, 1985. doi: 10.1080/01621459.1985.10478157. URL <http://www.tandfonline.com/doi/abs/10.1080/01621459.1985.10478157>.
- Carlos M. Carvalho, Jeffrey Chang, Joseph E. Lucas, Joseph R. Nevins, Quanli Wang, and Mike West. High-dimensional sparse factor modeling: Applications in gene expression genomics. *Journal of the American Statistical Association*, 103(484):1438–1456, 2008. ISSN 0162-1459. doi: 10.1198/016214508000000869.
- Merlise Clyde and Edward I George. Flexible empirical Bayes estimation for wavelets. *Journal of the Royal Statistical Society Series B*, 62(4):681–698, 2000. doi: 10.1111/1467-9868.00257.
- GTEX Consortium et al. The genotype-tissue expression (GTEx) pilot analysis: Multitissue gene regulation in humans. *Science*, 348(6235):648–660, 2015.
- Chris HQ Ding, Tao Li, and Michael I Jordan. Convex and semi-nonnegative matrix factorizations. *IEEE transactions on pattern analysis and machine intelligence*, 32(1):45–55, 2008.

- C Eckart and G Young. The approximation of one matrix by another of lower rank. *Psychometrika*, 1:211–218, 1936.
- Barbara E Engelhardt and Matthew Stephens. Analysis of population structure: a unifying framework and novel methods based on sparse factor analysis. *PLoS Genetics*, 6(9): e1001117, sep 2010.
- William Fithian, Rahul Mazumder, et al. Flexible low-rank statistical modeling with missing data and side information. *Statistical Science*, 33(2):238–260, 2018.
- J Kevin Ford, Robert C MacCallum, and Marianne Tait. The application of exploratory factor analysis in applied psychology: A critical review and analysis. *Personnel Psychology*, 39(2):291–314, 1986.
- Sylvia Frühwirth-Schnatter and Hedibert Freitas Lopes. Sparse Bayesian factor analysis when the number of factors is unknown. *arXiv preprint arXiv:1804.04231*, 2018.
- Chuan Gao, Christopher D Brown, and Barbara E Engelhardt. A latent factor model with a mixture of sparse and dense factors to model gene gene expression data with confounding effects. *arXiv:1310.4792v1*, 2013.
- Chuan Gao, Ian C. McDowell, Shiwen Zhao, Christopher D. Brown, and Barbara E. Engelhardt. Context specific and differential gene co-expression networks via Bayesian bi-clustering. *PLoS Computational Biology*, 12(7):1–39, 07 2016. doi: 10.1371/journal.pcbi.1004791. URL <https://doi.org/10.1371/journal.pcbi.1004791>.
- Zoubin Ghahramani and Matthew J Beal. Variational inference for Bayesian mixtures of factor analysers. In *Advances in neural information processing systems*, pages 449–455, 2000.
- Mark Girolami. A variational method for learning sparse and overcomplete representations. *Neural Computation*, 13(11):2517–2532, 2001.
- F Maxwell Harper and Joseph A Konstan. The Movielens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4):19, 2016.
- Trevor Hastie, Rahul Mazumder, Jason D Lee, and Reza Zadeh. Matrix completion and low-rank SVD via fast alternating least squares. *Journal of Machine Learning Research*, 16(1):3367–3402, 2015.
- Sepp Hochreiter, Ulrich Bodenhofer, Martin Heusel, Andreas Mayr, Andreas Mitterecker, Adetayo Kasim, Tatsiana Khamiakova, Suzy Van Sanden, Dan Lin, Willem Talloen, Luc Bijnens, Hinrich W. H. Göhlmann, Ziv Shkedy, and Djork-Arné Clevert. FABIA: factor analysis for bicluster acquisition. *Bioinformatics*, 26(12):1520–1527, 04 2010. ISSN 1367-4803. doi: 10.1093/bioinformatics/btq227. URL <https://doi.org/10.1093/bioinformatics/btq227>.
- Victoria Hore, Ana Viñuela, Alfonso Buil, Julian Knight, Mark I McCarthy, Kerrin Small, and Jonathan Marchini. Tensor decomposition for multiple-tissue gene expression experiments. *Nature Genetics*, 48(9):1094–1100, 2016.

- Tommi S. Jaakkola and Michael I. Jordan. Bayesian parameter estimation via variational methods. *Statistics and Computing*, 10:25–37, 2000.
- Iain M. Johnstone and Bernard W. Silverman. Empirical Bayes selection of wavelet thresholds. *The Annals of Statistics*, 33(4):1700–1752, 2005a.
- Iain M Johnstone and Bernard W Silverman. EBayesthresh: R and s-plus programs for empirical Bayes thresholding. *J. Statist. Soft*, 12:1–38, 2005b.
- Iain M Johnstone, Bernard W Silverman, et al. Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences. *The Annals of Statistics*, 32(4):1594–1649, 2004.
- Ian T Jolliffe, Nickolay T Trendafilov, and Mudassir Uddin. A modified principal component technique based on the lasso. *Journal of Computational and Graphical Statistics*, 12(3): 531–547, 2003.
- Julie Josse, Sylvain Sardy, and Stefan Wager. denoiseR: A package for low rank matrix estimation, 2018.
- Sylvia Kaufmann and Christian Schumacher. Identifying relevant and irrelevant variables in sparse factor models. *Journal of Applied Econometrics*, 32(6):1123–1144, 2017.
- Arto Klami. Poly-gamma augmentations for factor models. In *Asian Conference on Machine Learning*, pages 112–128, 2015.
- David Knowles and Zoubin Ghahramani. Nonparametric Bayesian sparse factor. *Annals of Applied Statistics*, 5(2B):1534–1552, 2011. doi: 10.1214/10-AOAS435.
- Roger Koenker and Ivan Mizera. Convex optimization, shape constraints, compound decisions, and empirical Bayes rules. *Journal of the American Statistical Association*, 109(506):674–685, 2014a.
- Roger Koenker and Ivan Mizera. Convex optimization in R. *Journal of Statistical Software*, 60(5):1–23, 2014b. doi: 10.18637/jss.v060.i05.
- Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. *SIAM review*, 51(3):455–500, 2009.
- D D Lee and H S Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755):788–791, 1999. doi: 10.1038/44565.
- Emma Lefrançois, Guadalupe Ortiz-muñoz, Axelle Caudrillier, Beñat Mallavia, Fengchun Liu, David M Sayah, Emily E Thornton, Mark B Headley, Tovo David, Shaun R Coughlin, Matthew F Krummel, Andrew D Leavitt, Emmanuelle Passegué, and Mark R Looney. The lung is a site of platelet biogenesis and a reservoir for haematopoietic progenitors. *Nature*, 544(7648):105–109, 2017. doi: 10.1038/nature21706.
- Yew Jin Lim and Yee Whye Teh. Variational Bayesian approach to movie rating prediction. In *Proceedings of KDD cup and workshop*, volume 7, pages 15–21. Citeseer, 2007.

- Vinicius Diniz Mayrink, Joseph Edward Lucas, et al. Sparse latent factor models with interactions: Analysis of gene expression data. *The Annals of Applied Statistics*, 7(2): 799–822, 2013.
- Rahul Mazumder, Trevor Hastie, and Robert Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11(Aug): 2287–2322, 2010.
- Shinichi Nakajima and Masashi Sugiyama. Theoretical analysis of Bayesian matrix factorization. *Journal of Machine Learning Research*, 12:2583–2648, 2011.
- Shinichi Nakajima, Masashi Sugiyama, S Derin Babacan, and Ryota Tomioka. Global analytic solution of fully-observed variational bayesian matrix factorization. *Journal of Machine Learning Research*, 14(Jan):1–37, 2013.
- Art B Owen and Jingshu Wang. Bi-cross-validation for factor analysis. *Statistical Science*, 31(1):119–139, 2016.
- Iosifina Pournara and Lorenz Wernisch. Factor analysis for gene regulatory networks and transcription factor activity profiles. *BMC bioinformatics*, 8:61, 2007. ISSN 1471-2105. doi: 10.1186/1471-2105-8-61.
- Tapani Raiko, Alexander Ilin, and Juha Karhunen. Principal component analysis for large scale problems with lots of missing values. In *European Conference on Machine Learning*, pages 691–698. Springer, 2007.
- Veronika Ročková and Edward I George. Fast Bayesian factor analysis via automatic rotations to sparsity. *Journal of the American Statistical Association*, 111(516):1608–1622, 2016.
- Donald B Rubin. Inference and missing data. *Biometrika*, 63(3):581–592, 1976.
- Donald B Rubin and Dorothy T Thayer. EM algorithms for ML factor analysis. *Psychometrika*, 47(1):69–76, 1982.
- Chiara Sabatti and Gareth M James. Bayesian sparse hidden components analysis for transcription regulation networks. *Bioinformatics*, 22(6):739–746, 2005.
- Matthias Seeger and Guillaume Bouchard. Fast variational Bayesian inference for non-conjugate matrix factorization models. In *Artificial Intelligence and Statistics*, pages 1012–1018, 2012.
- Sanvesh Srivastava, Barbara E Engelhardt, and David B Dunson. Expandable factor analysis. *Biometrika*, 104(3):649–663, 2017.
- Oliver Stegle, Leopold Parts, Richard Durbin, and John Winn. A Bayesian framework to account for complex non-genetic factors in gene expression levels greatly increases power in eqtl studies. *PLoS Computational Biology*, 6(5):e1000770, 2010.

Oliver Stegle, Leopold Parts, Matias Piipari, John Winn, and Richard Durbin. Using probabilistic estimation of expression residuals (PEER) to obtain increased power and interpretability of gene expression analyses. *Nature Protocols*, 7(3):500–507, 2012.

Genevieve L Stein-O’Brien, Raman Arora, Aedin C Culhane, Alexander V Favorov, Lana X Garmire, Casey S Greene, Loyal A Goff, Yifeng Li, Aloune Ngom, Michael F Ochs, et al. Enter the matrix: factorization uncovers knowledge from omics. *Trends in Genetics*, 34(10):790–805, 2018.

Matthew Stephens. False discovery rates: a new deal. *Biostatistics*, 18(2):275–294, Apr 2017. doi: 10.1093/biostatistics/kxw041.

D C Thomas, J Siemiatycki, R Dewar, J Robins, M Goldberg, and B G Armstrong. The problem of multiple inference in studies designed to generate hypotheses. *American Journal of Epidemiology*, 122(6):1080–95, Dec 1985.

Michael E Tipping. Sparse Bayesian learning and the relevance vector machine. *Journal of Machine Learning Research*, 1(Jun):211–244, 2001.

Michalis K. Titsias and Miguel Lázaro-Gredilla. Spike and slab variational inference for multi-task and multiple kernel learning. In J. Shawe-Taylor, R. S. Zemel, P. L. Bartlett, F. Pereira, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 24*, pages 2339–2347. Curran Associates, Inc., 2011. URL <http://papers.nips.cc/paper/4305-spike-and-slab-variational-inference-for-multi-task-and-multiple-kernel-learning.pdf>.

Wei Wang. *Applications of Adaptive Shrinkage in Multiple Statistical Problems*. PhD thesis, The University of Chicago, 2017.

Mike West. Bayesian factor regression models in the “large p, small n” paradigm. *Bayesian Statistics 7 - Proceedings of the Seventh Valencia International Meeting*, pages 723–732, 2003. ISSN 08966273. doi: 10.1.1.18.3036.

David P. Wipf and Srikantan S. Nagarajan. A new view of automatic relevance determination. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 1625–1632. Curran Associates, Inc., 2008. URL <http://papers.nips.cc/paper/3372-a-new-view-of-automatic-relevance-determination.pdf>.

Daniela M. Witten, Robert Tibshirani, and Trevor Hastie. A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis. *Biostatistics*, 10(3):515–534, 2009. doi: 10.1093/biostatistics/kxp008.

Zhengrong Xing, Peter Carbonetto, and Matthew Stephens. Smoothing via adaptive shrinkage (smash): denoising Poisson and heteroskedastic Gaussian signals. *arXiv preprint arXiv:1605.07787*, 2016.

Dan Yang, Zongming Ma, and Andreas Buja. A sparse singular value decomposition method for high-dimensional data. *Journal of Computational and Graphical Statistics*, 23(4):923–942, 2014.

Shiwen Zhao, Barbara E Engelhardt, Sayan Mukherjee, and David B Dunson. Fast moment estimation for generalized latent Dirichlet models. *Journal of the American Statistical Association*, pages 1–13, 2018.

Hui Zou, Trevor Hastie, and Robert Tibshirani. Sparse principal component analysis. *Journal of Computational and Graphical Statistics*, 15(2):265–286, 2006. ISSN 1061-8600. doi: 10.1198/106186006X113430.