

# Population Priors for Matrix Factorization

Anonymous authors  
Paper under double-blind review

Notation tracking.  
are in Blue.

The references and citations are not  
shown as ? in this document.

## Abstract

We develop an empirical Bayes prior for probabilistic matrix factorization. Matrix factorization models each cell of a matrix with two latent variables, one associated with the cell’s row and one associated with the cell’s column. How to set the priors of these two latent variables? Drawing from empirical Bayes principles, we consider estimating the priors from data, to find those that best match the populations of row and column latent vectors. Thus we develop the twin population prior. We develop a variational inference algorithm to simultaneously learn the empirical priors and approximate the corresponding posterior. We evaluate this approach with both synthetic and real-world data on diverse applications: movie ratings, book ratings, single-cell gene expression data, and musical preferences. Without needing to tune Bayesian hyperparameters, we find that the twin population prior leads to high-quality predictions, outperforming manually tuned priors.

## 1 Introduction

This paper is about empirical Bayes methods for setting the priors in Bayesian matrix factorization (??). Matrix factorization models each cell of a matrix with two latent variables, one associated with its row and one associated with its column. Matrix factorization has found broad applications across many fields, including studying consumer behavior, understanding legislative patterns, assessing pharmaceutical impacts, and exploring social networks (????).

Suppose  $X_{i,j}$  is the observed entry in row  $i$  and cell  $j$ , such as user  $i$ ’s rating of movie  $j$ . As a hierarchical model, a matrix factorization generates the data from the following process:

$$\mathbf{U}_i \sim P_{\theta^r}(\mathbf{U}) \quad (1)$$

$$\mathbf{V}_j \sim P_{\theta^c}(\mathbf{V}) \quad (2)$$

$$X_{i,j} \sim P(X_{i,j} | \mathbf{U}_i, \mathbf{V}_j). \quad (3)$$

Here  $\mathbf{U}_i$  and  $\mathbf{V}_j$  are per row and per column specific latent vectors, and  $P_{\theta^r}$  and  $P_{\theta^c}$  are the priors, each with hyperparameters  $\theta^r$  and  $\theta^c$ .

This formulation encompasses many factorization models. In Gaussian matrix factorization (?), the priors are Gaussians and  $X_{i,j}$  is drawn from a Gaussian with mean  $\mathbf{U}_i^\top \mathbf{V}_j$  and variance  $\sigma^2$ . In Poisson matrix factorization (???), the priors are over the positive reals and  $X_{i,j}$  is drawn from a Poisson with rate  $\mathbf{U}_i^\top \mathbf{V}_j$ .

An observed matrix of data  $\mathbf{X}$  then defines a posterior distribution  $P(\mathbf{U}, \mathbf{V} | \mathbf{X})$  over the row variables and the column variables. The posterior can provide interpretations of the data and an avenue to form predictions about missing entries, for example, for a recommendation system.

The prior distributions on the row variables and column variables significantly impact the quality of the model posterior. How should we set them? Practitioners typically assume a simple parametric family for the priors, such as a Gaussian or a Gamma, and then find the prior hyperparameters best suited to the data, e.g., with cross-validation (?). This approach can be effective, but it is expensive and only allows for priors from a simple parametric family.

In this paper, we develop an empirical Bayes (EB) methodology for setting the priors (??), learning them from data. The EB idea is to set the priors using the data, for example by finding the one that maximizes the marginal log-likelihood of the data. It is applicable for priors over variables that have repeated draws from them, such as the row and column variables in matrix factorization. EB has found applications in applied sciences in diverse fields such as astronomy (?), actuarial sciences (?), genomics (??), economics (??), and survey sampling (?). EB priors have been successfully employed in simple hierarchical models, such as in variational autoencoders (??).

Matrix factorization, however, provides a different type of application of EB. In matrix factorization, there are two priors to set, one for the row variables and one for the column variables, and the data that informs us about them is overlapping. Thus, we will find empirical Bayes priors for both row variables and column variables. The result is the *twin empirical Bayes prior* (TwinEB), a practical EB method for matrix factorization. Other methods for EB on matrix factorization include ? and ?; we discuss these below.

Specifically, we model the two priors with mixture distributions, one for each prior. We use mixtures since they are a flexible family of distributions that can approximate a wide range of distributions (??). We use variational inference (??) to simultaneously estimate the priors and approximate the corresponding posterior. We verify the efficiency and the robustness of this approach with real-world data about recommendation systems and computational biology, and for both Poisson matrix factorization and Gaussian matrix factorization.

We summarize the contributions as follows:

1. We develop the *twin population prior*, an EB prior for Bayesian matrix factorization.
2. We derive a variational inference algorithm to approximate the matrix factors' posteriors and learn the twin population prior simultaneously.
3. We study the twin population prior on both synthetic and real data and with two types of factorization. The automatically learned EB prior performs as well as the best prior chosen retrospectively.

## 2 Related work

One approach to setting the hyperparameters of priors involves using hierarchical Bayesian models (?). In this line of work, the prior's hyperparameters are treated as unknown and assigned a prior with hyperpriors to be determined. ? employs this method for introducing hierarchical Poisson factorization, and ? applies it to gene signature discovery from scRNAseq data. Our work involves more expressive priors (mixtures), which leads to a more flexible class of models while keeping a simpler model as no extra variables for the hyperparameters are introduced.

A fruitful direction of research for setting priors for matrix factorization has also been the use of empirical Bayes, a methodology by which one tries to find the prior that matches, in some sense, the population distribution of the data. In ?, the authors formulate prior elicitation for matrix factorization (MF) as steps in a variational expectation maximization algorithm, and in that context, find that it is equivalent to solving the EB normal means (EBNM) problem (?). Our proposed approach is closely related to ?, but we propose to update both priors simultaneously which bypasses potentially costly numerical integration required in the EBNM step. A related line of research makes the connection to EBNM via random-matrix theory, focusing specifically on denoising principal components (?). Closely related to TwinEB, ? optimizes hyperparameters to match the prior predictive distribution to statistics computed from the data, and derives closed-form solutions for PMF and its hierarchical extensions. In Section ??, we compare our method to ? and ?.

Empirical Bayes (EB) priors have been explored in the context of variational autoencoders (VAEs) (?) under the name aggregated posterior, average encoding distribution (?) or VampPrior (?). The VampPrior learns an amortized posterior and a prior over the latent variables using a shared neural network. It models a prior on the rows only and was used to address posterior collapse (?) or to learn disentangled representations (?). Our method focuses on matrix factorization, and we derive EB priors for both row and column latent variables.

### 3 Empirical Bayes priors for probabilistic matrix factorization

Our goal is to develop empirical Bayes (EB) priors for Bayesian matrix factorization models. We will focus here on Poisson matrix factorization (PMF). In the supplement, we derive EB priors for Gaussian matrix factorization (GMF).

With matrix factorization, the presence of repeated and identically distributed latent variables for each row and each column provides the opportunity to learn their prior distribution from data. This is a form of empirical Bayes (??) that prescribes a *population prior* (see ??).

This population prior aims to align the model’s marginal distribution of observations with the observed population distribution. In the special case of matrix factorization, there are two distinct populations: the populations of row vectors, and the population of column vectors. With TwinEB, we learn one prior for each population. This is a form of hierarchical modeling without introducing an extra layer of latent variables.

**Notations.** Define  $[N] := 1, \dots, N$ . Let  $\mathbf{X} = [X_{i,j}] \in \mathbb{N}^{N \times D}$  represent the measured feature  $j$  in individual

We employ boldface symbols to distinguish vectors from scalars. We recapitulate the notations below:

- $\mathbf{X} \in \mathbb{N}^{N \times D}$ : the observed data matrix;  $X_{i,j}$  is the entry in row  $i$ , where  $i \in [N]$  and  $j \in [D]$ . The dimension of the latent variables (and column  $j$ .
- $N$  the number of rows (e.g., individuals),  $D$  the number of factors) is denoted as columns (e.g., features).
- $P_{\text{row}}^*(\mathbf{X}_{i,:})$  is the unknown distribution of rows vectors, and  $P_{\text{col}}^*(\mathbf{X}_{:,j})$  the one of columns.
- $\mathbf{U}_i \in \mathbb{R}^L$  and  $\mathbf{V}_j \in \mathbb{R}^L$  are model latent variables for row  $i$  and column  $j$ ; of dimension  $L$ .
- $\pi$  is the prior distribution of the latent variables.

#### 3.1 Background: Population priors for simple hierarchical models

In this paper, we assume to have data  $\mathbf{X}$  that is drawn from an unknown distribution  $P^*$ , also called *population distribution*, that we would like to model with a probabilistic model consisting of latent variables  $\mathbf{Z}$ , a prior over the latent variables  $\pi$ , and a likelihood function  $P(\mathbf{X} | \mathbf{Z})$ .

A crucial step in Bayesian statistics is the choice of the prior distribution; if done arbitrarily, it can lead to suboptimal posterior inference (?). We choose to follow an empirical Bayes principle that prescribes a *population prior* (??). This prior, by design, aligns the model’s marginal distribution of observations with the population distribution  $P^*(\mathbf{X})$ .

We first focus on a family of latent variable models called simple hierarchical models (?). The joint distribution factorizes as follows:

$$P(\mathbf{Z}, \mathbf{X}) = \prod_{i=1}^N P_{\pi} \pi(z_i) P_{\theta}(\mathbf{x}_i | z_i), \quad (4)$$

where  $\theta$  and  $\mathbf{Z} = [z_i]$  are global and local latent variables respectively and observations  $\mathbf{X} = [\mathbf{x}_i]$  are exchangeable conditioned on the latent variables. The variable  $\pi$  indexes the  $\pi$  is the prior distribution of  $z_i$ . We assume that the global latent variables are fixed (e.g., they are learned through MAP estimation). We use a subscript notation to indicate fixed variables. To further To simplify notations, we then focus on the marginal likelihood of a single observation,  $\mathbf{x}$ , and its corresponding local latent variable  $z$ .

Let  $P^*(\mathbf{x})$  be the true (unobserved) distribution of observations. An empirical Bayes criterion is that the marginal distribution of observations under the model,  $P_{\pi, \theta}(\mathbf{x})$ , should align denoted as  $P_{\pi}(\mathbf{x})$ , should

match with their true population distribution  $P^*(\mathbf{x})$  (?), that is:

$$\begin{aligned} P^*(\mathbf{x}) &= P_{\pi, \theta \pi}(\mathbf{x}) \\ &= \int P_{\pi} \pi(z) P_{\theta}(\mathbf{x} | z) dz. \end{aligned} \quad (5)$$

Our goal is to set  $\pi$  such that, for a fixed  $\theta$ ,  $\pi$  such that Equation (??) holds. For a fixed likelihood function  $P_{\theta}(\mathbf{x} | z)$  (i.e.,  $\theta$  is fixed), the The expression for the prior  $P_{\pi^* \pi}$  that satisfies this conditions is:

$$\begin{aligned} P_{\pi^* \pi}(z) &\approx \int P_{\pi, \theta \pi}(z | \mathbf{x}) P^*(\mathbf{x}) d\mathbf{x} \\ &= \mathbb{E}_{P^*(\mathbf{x})} [P_{\pi, \theta}(z | \mathbf{x})]_{P^*(\mathbf{x})} [P_{\pi}(z | \mathbf{x})], \end{aligned} \quad (6)$$

where  $P_{\pi, \theta}(z | \mathbf{x}) P_{\pi}(z | \mathbf{x})$  is the (local) posterior distribution of the latent variable  $z$  given the observation  $\mathbf{x}$  under the model. The definition presents two issues: the unknown true marginal population distribution  $P^*(\mathbf{x})$ , and the fact that the target prior  $P_{\pi^* \pi}(z)$  is on both sides of Equation (??), explicitly on the left and implicitly via the posterior on the right. The research literature has approximated Equation (??) with Monte Carlo estimates of  $P^*(\mathbf{x})$  and variational inference of  $P(z | \mathbf{x}) P_{\pi}(z | \mathbf{x})$  (??).

### 3.2 Population priors for probabilistic matrix factorization

Our goal is to develop population priors for Bayesian matrix factorization models. The challenge is that unlike simple hierarchical models, there is no distinction between local and global random variables, rather latent variables denote row- and column-specific random variables.

#### 3.2.1 Twin population priors

We establish population priors for two latent variables, one for the rows  $P_{\mathbf{V}^*}(\mathbf{U}_i)$  latent variables of the rows  $\pi_{\text{row}}(\mathbf{U}_i)$  and one for the columns  $P_{\mathbf{U}^*}(\mathbf{V}_j)$  latent variables of the columns  $\pi_{\text{col}}(\mathbf{V}_j)$ . These priors will match two different populations, one of the row vectors and one of the column vectors:

$$\begin{aligned} P_{\text{row}}^*(\mathbf{X}_{i,:}) &:= \text{Population distribution of row vectors} \\ P_{\text{col}}^*(\mathbf{X}_{:,j}) &:= \text{Population distribution of column vectors} \end{aligned}$$

We recall that the populations distributions are unknowns distribution, from which rows are samples as we gather more observations and columns are samples as we gather more features.

We begin with population priors for row latent variables. As in section ??, we specify the prior based on an empirical Bayes principle such that the true marginal distribution of the rows  $P_{\text{row}}^*(\mathbf{X}_{i,:})$  is aligned with the distribution of the rows under the model  $P_{\mathbf{V}^*}(\mathbf{X}_{i,:})$ , that is:

$$P_{\text{row}}^*(\mathbf{X}_{i,:}) = P_{\mathbf{V}^*}(\cdot) \equiv \int P_{\pi_{\text{row}}}(\mathbf{U}_i) \prod_{j=1}^D P(X_{i,j} | \mathbf{U}_i, \mathbf{V}_j) d\mathbf{U}_i, \quad (7)$$

for a fixed set of column variables  $\mathbf{V}$ . For Equation (??) to hold, a population prior should be used:

$$P_{\mathbf{V}^*} \pi_{\text{row}}(\mathbf{U}_i; \mathbf{V}) = \mathbb{E}_{P_{\text{row}}^*(\mathbf{X}_{i,:})} [P_{\mathbf{V}^* \pi_{\text{row}}}(\mathbf{U}_i | \mathbf{X}_{i,:}, \mathbf{V})], \quad (8)$$

where we explicited the dependence of the EB prior on the column latent variables  $\mathbf{V}$ . Similarly for columns, for a fixed set of row latent variables  $\mathbf{U}$ , the empirical Bayes criterion is:

$$P_{\text{col}}^*(\mathbf{X}_{:,j}) = P_{\mathbf{U}^*}(\cdot) \equiv \int P(\mathbf{V}_i) \prod_{i=1}^N P_{\mathbf{U}_j}(X_{i,j} | \mathbf{U}_i, \mathbf{V}_j) d\mathbf{V}_{ij}. \quad (9)$$

The prior that satisfies this criterion is the column population prior:

$$\underline{P_U}^* \pi_{\text{col}}(\mathbf{V}_j; \mathbf{U}) = \mathbb{E}_{P_{\text{col}}^*(\mathbf{X}_{:,j})} [P_{\underline{U}} \pi_{\text{col}}(\mathbf{V}_j \mid \mathbf{X}_{:,j}, \mathbf{U})]. \quad (10)$$

Since there are two populations in need of prior specification, we call Equations (?? and ??) the *twin* population priors.

We have established the form of population priors in probabilistic matrix factorization. Next we focus on how to estimate the twin population priors and how to approximate posterior inference under them.

In the remainder of the paper, we focus on the Poisson matrix factorization (PMF). We derive the priors for Gaussian matrix factorization (GMF) in the supplement.

### 3.3 Twin EB prior for Poisson matrix factorization

In Poisson matrix factorization (PMF), the row and column latent variables  $\mathbf{U}_i$  and  $\mathbf{V}_j$  are non-negative  $L$ -vectors and the likelihood in Equation (??) is Poisson:

$$X_{i,j} \mid \mathbf{U}_i, \mathbf{V}_j \sim \text{Poisson}(\mathbf{U}_i^T \mathbf{V}_j), \quad (11)$$

for  $i \in [N], j \in [D]$ .

The log-likelihood of the data is:

$$\begin{aligned} \log P(\mathbf{X} \mid \mathbf{U}, \mathbf{V}) &= \log \prod_{i=1}^N \prod_{j=1}^D P(X_{i,j} \mid \mathbf{U}_i, \mathbf{V}_j) \\ &= \sum_{i,j} \log \text{Poisson}(X_{i,j}; \mathbf{U}_i^T \mathbf{V}_j). \end{aligned} \quad (12)$$

Some methods place Gamma priors on  $\mathbf{V}_j$  and  $\mathbf{U}_i$  (?). Note that this is a Bayesian formulation of non-negative matrix factorization (?).

To compute the population prior for the rows,  $P_{\mathbf{V}}(\mathbf{U}_i) = \mathbb{E}_{P_{\text{row}}^*(\mathbf{X}_{i,:})} [P_{\mathbf{V}}(\mathbf{U}_i \mid \mathbf{X}_{i,:})] \pi_{\text{row}}(\mathbf{U}_i) = \mathbb{E}_{P_{\text{row}}^*(\mathbf{X}_{i,:})} [P_{\mathbf{V}}(\mathbf{U}_i \mid \mathbf{X}_{i,:})]$ , we face two problems. Namely, we do not know the true population distribution of the rows  $P_{\text{row}}^*(\mathbf{X}_{i,:})$ , and the population prior  $P_{\mathbf{V}}(\mathbf{U}_i)$  appears on both sides of the equality.

To find the population prior, we first notice that a Monte Carlo estimate of ?? writes as:

$$\underline{P_V} \pi_{\text{row}}(\mathbf{U}_i; \mathbf{V}) \approx \frac{1}{N} \sum_{i'=1}^N P_{\mathbf{V}}(\mathbf{U}_i \mid \mathbf{X}_{i',:}, \mathbf{V}). \quad (13)$$

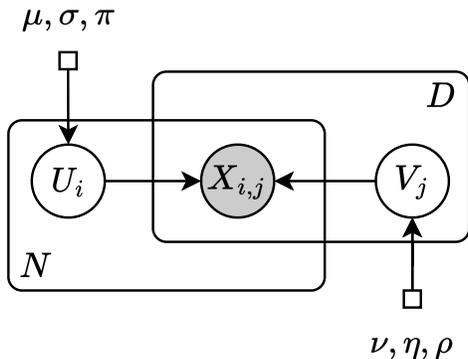
When the prior satisfies Equation (??), this property is called self-consistency (?).

The structure of Equations (??) suggests to use families of mixtures of parametric distributions to approximate the row and column population priors (?). Mixtures can approximate complex distributions when their number of components increases while having the convenience of remaining parametric (?). We choose to model the priors as, with as,

$$\underline{P_V} \pi_{\text{row}}(\mathbf{U}_i) := P_{\underline{\theta}^r} \pi_{\text{row}}(\mathbf{U}_i) := \sum_{k=1}^{K_r} \pi_{\omega_k} P_{\mu_k, \sigma_k}(\mathbf{U}_i), \quad (14)$$

$$\underline{P_U} \pi_{\text{col}}(\mathbf{V}_j) := P_{\underline{\theta}^c} \pi_{\text{col}}(\mathbf{V}_j) := \sum_{k=1}^{K_c} \rho_k P_{\nu_k, \eta_k}(\mathbf{V}_j), \quad (15)$$

where  $K^r$  and  $K^c$  are the number of components in the mixtures,  ~~$\underline{\theta}^r = \{\mu, \sigma, \pi\}$  and  $\underline{\theta}^c = \{\nu, \eta, \rho\}$~~   $\pi_{\text{row}} = \{\mu, \sigma, \omega\}$  and  $\pi_{\text{col}} = \{\nu, \eta, \rho\}$ . The locations  $\mu \in \mathbb{R}^{K_r \times L}$  and  $\nu \in \mathbb{R}^{K_c \times L}$ , the scales



**Figure 1: Twin population priors for Poisson matrix factorization model.** Shaded nodes are observed while other nodes represent latent random variables. The empty squares indicate that we will fit these priors to the data.

$\sigma \in \mathbb{R}^{K_r \times L}$  and  $\eta \in \mathbb{R}^{K_c \times L}$ , and the mixture weights  $\pi \in \Delta^{K_r}$  and  $\omega \in \Delta^{K_c}$  are the parameters of the mixtures priors. As  $K_r$  and  $K_c$  increase, the priors become more and more expressive (see Figure ??). Figure ?? shows a graphical model representation of matrix factorization with EB priors.

In a classical empirical Bayes setup, the idea is to set the priors that maximize the marginal likelihood of the data:

$$\hat{\theta}_{\text{row}}, \hat{\theta}_{\text{col}} = \arg \max_{\theta^r, \theta^c} \log P(\mathbf{X}; \theta^r \theta_{\text{row}}, \theta^c \theta_{\text{col}}), \quad (16)$$

where

$$\log P(\mathbf{X}; \theta^r \theta_{\text{row}}, \theta^c \theta_{\text{col}}) = \log \int P_{\theta^r \theta_{\text{row}}}(U) P_{\theta^c \theta_{\text{col}}}(V) P(\mathbf{X} | U, V) dU dV. \quad (17)$$

In ??, we find  $\theta^r, \theta^c$  at the same time as we approximate the model posterior  $P(U, V | \mathbf{X}; \theta^r, \theta^c)$ .

### 3.4 Posterior inference in PMF with twin EB priors

Given data  $\mathbf{X}$ , our goal is to calculate the posterior  $P(U, V | \mathbf{X}; \theta^r, \theta^c)$ , which also depends on our choice of priors  $\theta^r, \theta^c$ . The challenges are that this posterior is intractable (for any prior) and we simultaneously want to fit the priors to satisfy the EB criterion in ??.

Our strategy will be as follows. We will use variational inference (VI) (?) to approximate the posterior, taking gradient steps in the variational objective with respect to the posterior approximation (the variational family). At the same time, however, the variational objective of VI is an approximation (lower bound) of the log-marginal from ?. So we also take gradient steps with respect to the EB priors to maximize it. The result is an algorithm that simultaneously approximates the posterior and learns the EB prior.

**The variational posterior.** Consider a parameterized mean-field variational family,

$$q_{\Lambda}(U, V | \mathbf{X}) = \prod_{i,l} q_{\lambda_{i,l}^r}(U_{i,l}) \prod_{j,l} q_{\lambda_{j,l}^c}(V_{j,l}), \quad (18)$$

This family has parameters for each row's latent vector and each column's latent vector,  $\lambda_i^r$  and  $\lambda_j^c$  respectively. We further define  $\Lambda^r := [\lambda_{i,l}^r]$ , and  $\Lambda^c := [\lambda_{j,l}^c]$ . The full set of variational parameters is  $\Lambda = \{\Lambda^r, \Lambda^c\}$ .

From the perspective of posterior inference, our goal is to set  $q_{\Lambda}$  to minimize the KL divergence to the exact posterior:

$$\hat{\Lambda} = \arg \min_{\Lambda} \text{KL}(q_{\Lambda}; P(\mathbf{U}, \mathbf{V} \mid \mathbf{X}; \underline{\theta}^r \theta_{\text{row}}, \underline{\theta}^c \theta_{\text{col}})). \quad (19)$$

In detail, the variational family is a bank of Log-Normals:

$$\lambda_{i,l}^r := (a'_{i,l}, b'_{i,l}), \quad (20)$$

$$\lambda_{j,l}^c := (a_{j,l}, b_{j,l}), \quad (21)$$

$$q_{\lambda_{i,l}^r}(U_{i,l}) := \mathcal{LN}(a'_{i,l}, b'_{i,l}), \quad (22)$$

$$q_{\lambda_{j,l}^c}(V_{j,l}) := \mathcal{LN}(a_{j,l}, b_{j,l}). \quad (23)$$

Each Log-Normal is parameterized by its natural parameters  $a$  and  $b$ :

$$\mathcal{LN}(x; a, b) \propto \exp\left(-\frac{a}{2b} \log(x) - \frac{(\log x)^2}{2b}\right).$$

To minimize the KL divergence in ??, VI optimizes the variational parameters  $\Lambda$  to, equivalently, maximize the evidence lower bound (ELBO) (?):

$$\begin{aligned} \mathcal{L}(\mathbf{X}; \Lambda, \underline{\theta}^r \theta_{\text{row}}, \underline{\theta}^c \theta_{\text{col}}) &= \mathbb{E}_{q_{\Lambda}(\mathbf{U}, \mathbf{V} \mid \mathbf{X})}[\log P(\mathbf{X} \mid \mathbf{U}, \mathbf{V})] \\ &+ \mathbb{E}_{q_{\Lambda}}[\log P(\mathbf{U}; \underline{\theta}^r, \underline{\theta}^c) + \log P(\mathbf{V}; \underline{\theta}^r, \underline{\theta}^c)]_{q_{\Lambda}} [\log P(\mathbf{U}; \theta_{\text{row}}, \theta_{\text{col}}) + \log P(\mathbf{V}; \theta_{\text{row}}, \theta_{\text{col}})] \\ &- \mathbb{E}_{q_{\Lambda}}[\log q_{\Lambda}(\mathbf{U}, \mathbf{V} \mid \mathbf{X})]. \end{aligned} \quad (24)$$

Here, we use gradient ascent to maximize  $\mathcal{L}(\mathbf{X}; \theta, \Lambda)$  with respect to  $\Lambda$  (?). We further use stochastic reparameterization gradients to take such steps (??).

**Maximum marginal likelihood.** At the same time, we would like to set the prior parameters to maximize the marginal likelihood of the data (??). The variational objective in ?? conveniently also provides a lower-bound on the marginal likelihood (?):

$$\log P(\mathbf{X}; \underline{\theta}^r \theta_{\text{row}}, \underline{\theta}^c \theta_{\text{col}}) \geq \mathcal{L}(\mathbf{X}; \theta, \Lambda). \quad (25)$$

So, we will also follow stochastic gradients of the ELBO with respect to the prior parameters  $\underline{\theta}^r, \underline{\theta}^c \theta_{\text{row}}, \theta_{\text{col}}$  to maximize  $\mathcal{L}(\mathbf{X}; \theta, \Lambda)$  with respect to  $\underline{\theta}^r, \underline{\theta}^c \theta_{\text{row}}, \theta_{\text{col}}$ . This strategy has been used in the context of linear regression (?).

**Twin EB.** Putting these two pieces together, our algorithm is a stochastic gradient ascent of the ELBO with respect to two sets of parameters. In optimizing with respect to  $\Lambda$ , we minimize the KL divergence between  $q_{\Lambda}$  and the posterior; in optimizing with respect to  $\underline{\theta}^r, \underline{\theta}^c \theta_{\text{row}}, \theta_{\text{col}}$ , we maximize the (approximate) marginal likelihood of the data.

We use the Adam algorithm for stochastic optimization (?) with a batch size of 128, and using ten particles to obtain unbiased noisy estimates of the gradient of the ELBO via the reparameterization trick (the particles are samples from  $q_{\Lambda}$  to estimate the expectation  $\mathbb{E}_{q_{\Lambda}}$  of the ELBO with Monte-Carlo).

The details of the algorithm are in Algorithm ??. Our implementation is available at [git@github.com:xxx/xxx](https://github.com/xxx/xxx).

**Algorithm 1** Variational inference for Poisson matrix factorization with twin EB priors

---

**Input:** Data  $\mathbf{X}$ , number of particles  $S$ , learning rate  $\zeta$ , number of iterations  $T$ , number of components  $K_r, K_c$ , number of latent dimensions  $L$ .

**Output:** Variational posterior parameters  $\Lambda^*$ , prior parameters  $\theta^{r*}, \theta^{c*}, \theta_{\text{row}}^*, \theta_{\text{col}}^*$ .

**Initialize:**  $\Lambda^{(0)}, \theta^{r(0)}, \theta^{c(0)}, \theta_{\text{row}}^{(0)}, \theta_{\text{col}}^{(0)}$ .

**for**  $t = 1$  to  $T$  **do**

**for**  $i = 1$  to  $N$ ,  $l = 1$  to  $L$  **do**

**for**  $s = 1$  to  $S$  **do**

      Sample  $\epsilon_{i,l}^{(s)} \sim \mathcal{N}(0, 1)$ .

      Compute  $U_{i,l}^{(s)} = \exp(a'_{i,l}{}^{(t-1)} + b'_{i,l}{}^{(t-1)} \epsilon_{i,l}^{(s)})$ .

**end for**

**end for**

**for**  $j = 1$  to  $D$ ,  $l = 1$  to  $L$  **do**

**for**  $s = 1$  to  $S$  **do**

      Sample  $\phi_{j,l}^{(s)} \sim \mathcal{N}(0, 1)$ .

      Compute  $V_{j,l}^{(s)} = \exp(a_{j,l}{}^{(t-1)} + b_{j,l}{}^{(t-1)} \phi_{j,l}^{(s)})$ .

**end for**

**end for**

  Estimate  $\mathcal{L}(\mathbf{X}; \theta^{r^{t-1}}, \theta^{c^{t-1}}, \Lambda^{(t-1)}) - \mathcal{L}(\mathbf{X}; \theta_{\text{row}}^{t-1}, \theta_{\text{col}}^{(t-1)}, \Lambda^{(t-1)})$  using Monte-Carlo in Equation (??) with samples  $\mathbf{U}^{(s)}, \mathbf{V}^{(s)}$  in place of  $\mathbb{E}_{q_\Lambda}$ .

$\Lambda^{(t)}, \theta^{r(t)}, \theta^{c(t)} \leftarrow \text{Adam}(\nabla_{(\theta^r, \theta^c, \Lambda)} \mathcal{L}, \zeta) \Lambda^{(t)}, \theta_{\text{row}}^{(t)}, \theta_{\text{col}}^{(t)} \leftarrow \text{Adam}(\nabla_{(\theta_{\text{row}}, \theta_{\text{col}}, \Lambda)} \mathcal{L}, \zeta)$ .

**end for**

**return**  $\Lambda^{(T)}, \theta^{r(T)}, \theta^{c(T)}, \theta_{\text{row}}^{(T)}, \theta_{\text{col}}^{(T)}$ .

---

## 4 Experiments

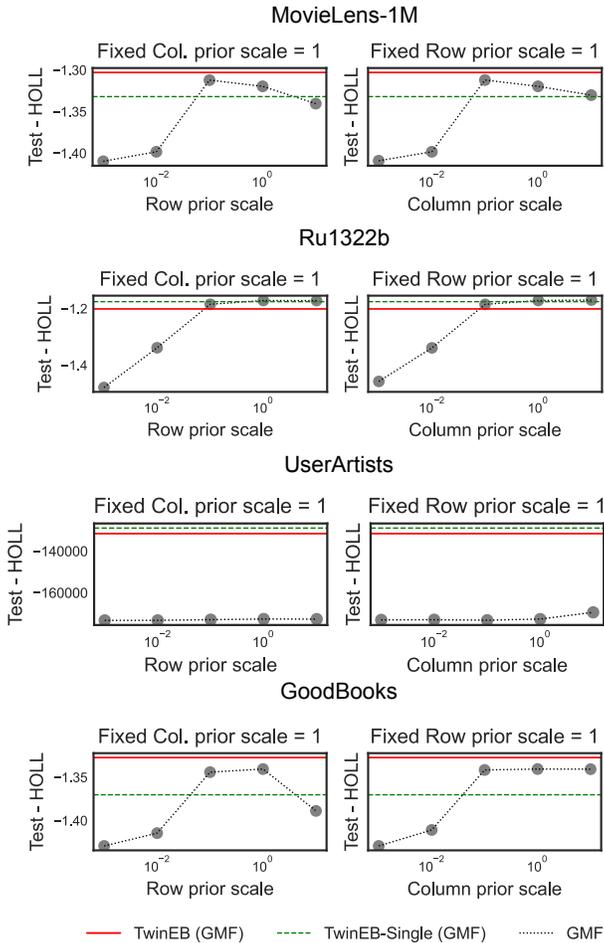
We studied Algorithm ?? in several real-world matrix factorization settings: book ratings, movie ratings, artist preferences, and single-cell RNA sequence gene counts. For all datasets, we studied both Gaussian and Poisson matrix factorization. In all datasets, we set  $L = 15$ . We found that TwinEB performs as well or better than manually searching for a parameterized prior, and performs as well or better than setting a simple parameterized prior by empirical Bayes.

### 4.1 Datasets

In this section we first review four real-world datasets, ranging from user-preferences to genomics, and then explore the impact of twin population priors on the performance of Gaussian and Poisson matrix factorization. In our experiments, we fix the number of row and column mixtures to  $K_r = 70$  and  $K_c = 100$  respectively.

**MovieLens 1M.** This dataset comprises 1 million ratings from 6,000 users (rows) on 4,000 movies (columns) (?). The ratings are on a scale of 1 to 5. Sparsity of this dataset, defined as the number of nonzero elements divided by the total number of entries, is 0.04. We can use matrix factorization to capture different aspects of user preferences and movie characteristics. Specifically,  $U_{i,l}$  may signify user  $i$ 's affinity for aspect  $l$  (e.g., genre), while  $V_{j,l}$  may represent the degree to which movie  $j$  exhibits aspect  $l$ .

**Ru1322b.** We analyze single cell gene expression data from a patient with small cell lung cancer (?). This dataset comprises 4,000 highly variable genes (columns) across 5,308 cells (rows). For the GMF family, we applied a two-step transformation. First, we performed a log transformation on the counts after adding a pseudo-count of one. Then, we standardized the non-zero elements. Sparsity of this dataset is 0.16. Each entry of the matrix denotes the number of transcripts of gene  $j$  in cell  $i$ . We explain the gene expression matrix via  $L$  gene-modules, with  $U_{i,l}$  as the activity of module  $l$  in cell  $i$ , and  $V_{j,l}$  as gene  $j$ 's contribution to module  $l$ . Matrix factorization can be used as exploratory data analysis (finding gene modules associated with malignancy in cancer) or as a component in a more complex analysis (causal inference (?)). See the supplement for details on preprocessing of the sequencing data.



**Figure 2: Twin population priors induce robustness to prior selection** The held-out likelihood is sensitive to the choice of the prior hyper-parameters. GMF endowed with population priors on both row and column latent variables, TwinEB (GMF), achieves comparable or better results than other methods. Each sub-panel displays held-out log-likelihood from adjusting the column prior variance with a fixed row prior, while the right sub-panel does the opposite, varying the row prior variance with a constant column prior. We demonstrate four datasets, from the top to bottom, MovieLens 1M, Ru1322b-scRNAseq, UserArtists, and GoodBooks. In all datasets, we set  $L = 15$ . Similar results hold for other values of  $L$  (see supplement).

**UserArtists** We use the data introduced in ?, comprising 92,834 user-listened artist relations, across 1,892 users and 17,632 artists, with a maximum value of 352,698. Similar to MovieLens 1M, we use matrix factorization to explain different aspects of user preferences and artist groupings. Sparsity of this dataset is 0.003. We note that this dataset was analyzed in ?.

**GoodBooks** Contains 6 million ratings across 51,288 users and 10,000 books<sup>1</sup>. Ratings range from 1 to 5. Sparsity of this dataset is 0.01.

#### 4.2 Evaluation Metric and Baselines

We evaluate each model based on the likelihood of held-out data. For a given row-wise data split into held-in and held-out rows, let  $\#\{i\}$  denote the set of  $N^{\text{out}}$  masked entries of the held-out rows. We estimate the

<sup>1</sup>Accessed at <https://github.com/zygmuntz/goodbooks-10k>

held-out log-likelihood as follows:

$$\frac{1}{N^{\text{out}}} \sum_{i=1}^{N^{\text{out}}} \log P(i | X^{\text{in}}) \approx \tag{26}$$

$$\frac{1}{N^{\text{out}}} \sum_{i=1}^{N^{\text{out}}} \log \frac{1}{M} \sum_{m=1}^M P(i | \mathbf{U}^{(m)}, \mathbf{V}^{(m)}), \tag{27}$$

where  $M$  is the number of Monte Carlo samples from  $q_{\Lambda}(\mathbf{U}, \mathbf{V})$ .

We randomly assign 20% of the rows as the test set, and the rest as training data. We then mask 20% of the entries at random, and train the model on this train set using ten random restarts. We use the 20% masked entries as a *validation set*. At test time, we put aside 30% of the entries of the test rows at random - these entries constitute the *test set* - then we train the model on 40% of the rest of the entries. This procedure measures *strong generalization* (?).

To compute the held-out likelihood, after training the model on the held-in data, we fix its column parameters and then learn row parameters of the held-out rows. We then report the likelihood of the masked entries in the test set. For each model, we report the test HOLL of the random restart that achieved the best validation HOLL. In our experiments, we set  $M = 500$ . See the supplement for the derivation of Equation (??) and more details on the experiments.

**Baseline.** We evaluate the performance of the PMF (GMF) model with TwinEB against multiple baselines, namely (i) TwinEB-Single, (ii) PMF (GMF). The former is a simple form of TwinEB where all latent dimensions have an identical prior, which we learn. The latter is a prior of the same family as the TwinEB, but with fixed hyperparameters. We compare TwinEB against a large choice of fixed parameters, akin to hyperparameter selection.

### 4.3 Results: Gaussian Matrix Factorization

We preprocess the data as follows. We standardize each column by subtracting the mean from non-zero entries and dividing the result by their standard deviation. We study two scenarios: (i) maintaining a fixed prior on row-wise variables while varying the prior on column-wise variables, and (ii) holding the prior on column-wise variables constant and adjusting the prior on row-wise variables. We set the variational family as well as the mixture components for the population priors to be Gaussian. We set the fixed prior to  $\mathcal{N}(1.0, 1.0)$  and vary the variance of the non-fixed one over  $\{0.001, 0.01, 0.1, 1.0, 10\}$ .

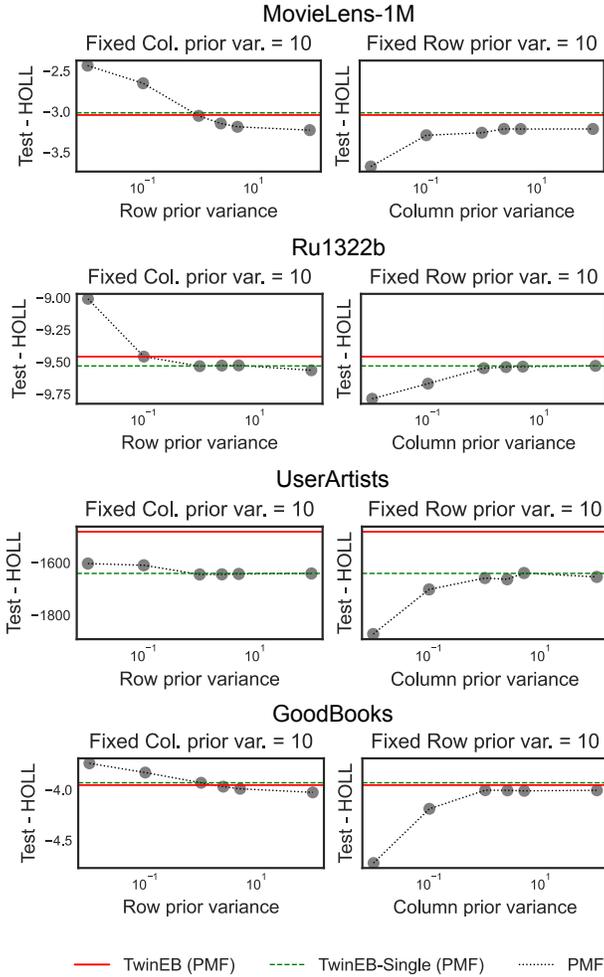
We compared the GMF model to TwinEB (GMF) and TwinEB-Single (GMF). We treat zero entries as missing values.

Figure ?? displays the outcomes for four real-world datasets. In GoodBooks and MovieLens-1M datasets, TwinEB achieves the highest test HOLL. For the UserArtists dataset, TwinEB does better than the fixed prior. Similar results were obtained by varying  $L$  to other values and are reported in Figures ??, ?? in ??.

We also compare to the method of ?, we used the corresponding package flashr, that currently supports GMF. We designed an imputation experiment where we held-out 10% of entries of a standardized matrix, and compare reconstruction accuracy on non-zero entries. We applied flashr to the above four real-world datasets. flashr crashes on all but the least sparse dataset, Ru1322b. For the Ru1322b dataset it achieves a mean absolute error of 0.64 vs 0.51 for the TwinEB. We note that it is not straightforward to compare our method to that of ?; the software implementation does not immediately support missing data and imputation.

### 4.4 Results: Poisson Matrix Factorization

For the MovieLens 1M and GoodBooks datasets, we binarize the matrix, setting entries to 'one' if a user has rated a movie and to 'zero' otherwise. For the Ru1322b dataset, we normalize the rows such that the sum of all rows are equal; we then round each value to the nearest integer. This is to account for the effect of library



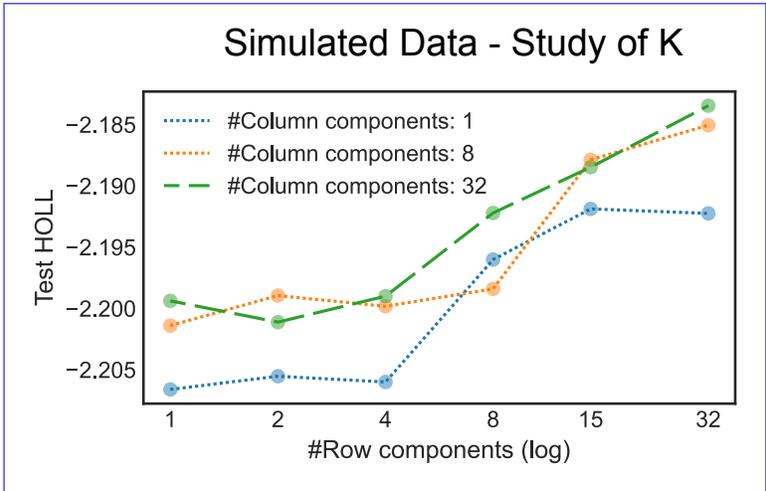
**Figure 3: Twin population priors induce robustness to prior selection.** As in Figure ??, but for PMF.

size (?). We treat zeros as missing data. In all models, we set the variational family to be Log-Normal, and the prior (or the mixture components for the population priors) to be Gamma. Similar to GMF, we vary the row or column prior parameters, while keeping the other one fixed. We parameterize the Gamma prior by its mean and variance  $\text{Gamma}(\mu, \sigma^2)$ , where  $\mu = \alpha/\beta$  and  $\sigma^2 = \alpha/\beta^2$ . In each scenario, we set the fixed prior to  $\text{Gamma}(1, 10)$ . For the varying prior, we set its mean to  $\mu = 1$  and change its variance along  $\{0.01, 0.1, 0.25, 1.0, 10.0, 100.0\}$ . Figure ?? shows the results.

In datasets except for GoodBooks, TwinEB outperforms TwinEBSingle. In the UserArtists dataset, TwinEB outperforms other methods. TwinEB, on average, takes about 1.5X the runtime of PMF (ranging from 1.1X to 2.1X). We point out that finding a good prior using grid-search incurs the sum of the cost of evaluations of the individual points in the grid, over 6.5X the running time of the PMF with TwinEB prior. We note that using a grid to find fixed priors (or alternatively, learning the scalar prior) does not guarantee reaching a best test HOLL. Indeed in the UserArtists dataset, TwinEB yields the best test HOLL.

We also compared TwinEB to the method of ?; given the data, it estimates values for the shape and rate parameters of the Gamma prior for both the row and column r.v.s. Table ?? displays the estimated hyperparameters and the resulting test HOLL.

PMF equipped with fixed priors values calculated from this methods yields Test HOLL that are on par or slightly worse than the TwinEB method. This method yielded negative hyperparameter estimates for the



**Figure 4: Increasing the number of mixtures improves the model performance.** Each line shows the average test-HOLL over 10 seeds for a fixed value for  $K_c$ , the number of column mixtures, and varying values of  $K_r$ , the number of row mixtures. Here, we set  $L = 64$  to its true simulated value. The results are similar for  $L = 32$  (see supplement).

MovieLens-1M and GoodBooks datasets. These parameters are effectively not usable, since the parameters of a Gamma must be positive. Hence, we emphasize that unlike ours, this methods cannot be used on all datasets.

**Table 1:** Test Heldout Loglikelihood from running PMF with prior hyperparameters sets using the method of ?. For MovieLens-1M and GoodBooks datasets, ? yielded inadmissible hyperparameters.

| Dataset      | test HOLL |               |
|--------------|-----------|---------------|
|              | ?         | TwinEB        |
| Ru1322b      | -9.53     | <b>-9.46</b>  |
| UserArtists  | -1,583    | <b>-1,481</b> |
| MovieLens-1M | NA        | <b>-3.04</b>  |
| GoodBooks    | NA        | <b>-3.95</b>  |

#### 4.5 Simulation: Complexity of the Prior

Here, we examine the performance of the population priors as the number of mixture components are varied. To this end, we simulate a 1,000 by 1,500 dataset, with  $L = 64$  dimensional row and column-wise r.v.s. We sample the row- and column-wise r.v.s from a mixture of 15 and 20 Gamma distributions respectively, the rate and shape parameters of which are sampled from a  $\text{Gamma}(1, 1)$ . The mixture weights are sampled from a  $\text{Dirichlet}(e_0, \dots, e_0)$  where the concentration parameter is  $e_0 = 10$ .

**Performance Increase with Number of Mixture Components.** The more mixture components, the more flexible the EB prior is. If learning a EB prior is beneficial, then we expect the performance to increase with the number of mixture components. We run PMF with population priors with varying number of row and column mixtures, setting the dimension of the latent variable to its true value, and, to avoid model misspecification, set the mixture prior family to Gamma. We then report the average of the test HOLL across ten different seeds. Figure ??, the log-likelihood of test held-out data increases with the number of mixture components in the row prior.

## 5 Discussion

We introduced the twin population priors for probabilistic matrix factorization. We derived a method to estimate the corresponding posterior using Monte Carlo and variational inference. On real-world data, this method finds a prior as good as the best parametric prior chosen retrospectively.

One area of further work is to extend this algorithm to tensor factorization (??). While in matrix factorization, each entry of the observed matrix is explained via two latent variables, tensor factorization models will involve more. One detail to address is how to formally define the population distribution associated with each.

### Broader Impact Statement

The authors do not foresee any negative impact of this work.

### References

- [Abhinav Agrawal and Justin Domke. Amortized Variational Inference for Simple Hierarchical Models. \*NeurIPS\*, 34:21388–21399, 2021.](#)
- [Joshua D Angrist, Peter D Hull, Parag A Pathak, and Christopher R Walters. Leveraging lotteries for school value-added: Testing and estimation. \*QJE\*, 132\(2\):871–919, 2017.](#)
- [David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. \*JASA\*, 112\(518\):859–877, 2017.](#)
- [Jo Bovy, David W. Hogg, and Sam T. Roweis. Extreme deconvolution: Inferring complete distribution functions from noisy, heterogeneous and incomplete observations. \*AOAS\*, 5\(2B\):1657–1677, 2011.](#)
- [Hans Bühlmann and Alois Gisler. \*A course in credibility theory and its applications\*, volume 317. Springer, 2005.](#)
- [John Canny. GaP: a factor model for discrete data. In \*SIGIR\*, pp. 122–129, 2004.](#)
- [Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. Second workshop on information heterogeneity and fusion in recommender systems \(hetrec2011\). In \*ACM RecSys\*, pp. 387–388, 2011.](#)
- [Ali Taylan Cemgil. Bayesian inference for nonnegative matrix factorisation models. \*COMPUT INTEL NEUROSC\*, 2009.](#)
- [Joseph M Chan, Álvaro Quintanal-Villalonga, Vianne Ran Gao, Yubin Xie, Viola Allaj, Ojasvi Chaudhary, Ignas Masilionis, Jacklynn Egger, Andrew Chow, Thomas Walle, et al. Signatures of plasticity, metastasis, and immunosuppression in an atlas of human small cell lung cancer. \*Cancer Cell\*, 39\(11\):1479–1496, 2021.](#)
- [Eliezer de Souza da Silva, Tomasz Kuł, Marcelo Hartmann, Arto Klami, et al. Prior specification for Bayesian matrix factorization via prior predictive matching. \*JMLR\*, 24\(67\):1–51, 2023.](#)
- [David B Dunson and Amy H Herring. Bayesian latent variable models for mixed discrete outcomes. \*Biostatistics\*, 6\(1\):11–25, 2005.](#)
- [Bradley Efron. \*Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction\*, volume 1. Cambridge University Press, 2012.](#)
- [Peter A Frost and James E Savarino. An empirical Bayes approach to efficient portfolio selection. \*JFQA\*, 21\(3\):293–305, 1986.](#)
- [Andrew Gelman. Prior distributions for variance parameters in hierarchical models \(comment on article by Browne and Draper\). 2006.](#)

- [Sean M Gerrish and David M Blei. Predicting legislative roll calls from text. In \*ICML\*, 2011.](#)
- [Prem Gopalan, Jake M Hofman, and David M Blei. Scalable recommendation with hierarchical Poisson factorization. In \*UAI\*, pp. 326–335, 2015.](#)
- [F Maxwell Harper and Joseph A Konstan. The movielens datasets: History and context. \*TiiS\*, 5\(4\):1–19, 2015.](#)
- [Lukas Heumos, Anna C. Schaar, et al. Best practices for single-cell analysis across modalities. \*Nat. Rev. Genet.\*, 24\(8\):550–572, 2023. ISSN 1471-0064.](#)
- [Matthew D Hoffman and Matthew J Johnson. ELBO surgery: yet another way to carve up the variational evidence lower bound. 1\(2\), 2016.](#)
- [Nikolaos Ignatiadis and Stefan Wager. Confidence intervals for nonparametric empirical Bayes analysis. \*JASA\*, 117\(539\):1149–1166, 2022.](#)
- [Mohsen Jamali and Martin Ester. A matrix factorization technique with trust propagation for recommendation in social networks. In \*RecSys\*, pp. 135–142, 2010.](#)
- [Wenhua Jiang and Cun-Hui Zhang. General maximum likelihood empirical Bayes estimation of normal means. \*Ann. Stat.\*, 2009.](#)
- [Hyunjik Kim and Andriy Mnih. Disentangling by factorising. In \*ICML\*, pp. 2649–2658. PMLR, 2018.](#)
- [Diederik P Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. \*arXiv:1412.6980\*, 2014.](#)
- [Diederik P Kingma and Max Welling. Auto-Encoding variational Bayes. \*arXiv:1312.6114\*, 2013.](#)
- [Tamara G Kolda and Brett W Bader. Tensor decompositions and applications. \*SIAM review\*, 51\(3\):455–500, 2009.](#)
- [Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. \*Computer\*, 42\(8\):30–37, 2009.](#)
- [Nan Laird. Nonparametric maximum likelihood estimation of a mixing distribution. \*JASA\*, 73\(364\):805–811, 1978.](#)
- [Hanna Mendes Levitin, Jinzhou Yuan, Yim Ling Cheng, Francisco JR Ruiz, Erin C Bush, Jeffrey N Bruce, Peter Canoll, Antonio Iavarone, Anna Lasorella, David M Blei, et al. De novo gene signature identification from single-cell RNA-seq with hierarchical Poisson factorization. \*Molecular systems biology\*, 15\(2\):e8557, 2019.](#)
- [Michael I Love, Wolfgang Huber, and Simon Anders. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. \*Genome biology\*, 15\(12\):1–21, 2014.](#)
- [Andriy Mnih and Russ R Salakhutdinov. Probabilistic matrix factorization. \*NeurIPS\*, 20, 2007.](#)
- [Sumit Mukherjee, Bodhisattva Sen, and Subhabrata Sen. A mean field approach to empirical Bayes estimation in high-dimensional linear regression. \*arXiv:2309.16843\*, 2023.](#)
- [T Tin Nguyen, Hien D Nguyen, Faicel Chamroukhi, and Geoffrey J McLachlan. Approximation by finite mixtures of continuous density functions that vanish at infinity. \*Cogent Mathematics & Statistics\*, 7\(1\):1750861, 2020.](#)
- [Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. \*NeurIPS\*, 32, 2019.](#)
- [Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In \*AISTATS\*, pp. 814–822. PMLR, 2014.](#)

- [John NK Rao and Isabel Molina. \*Small area estimation\*. John Wiley & Sons, 2015.](#)
- [Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In \*ICML\*, pp. 1278–1286. PMLR, 2014.](#)
- [Herbert E Robbins. \*An empirical Bayes approach to statistics\*. Springer, 1992.](#)
- [Ruslan Salakhutdinov and Andriy Mnih. Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In \*ICML\*, pp. 880–887, 2008.](#)
- [Aaron Schein, John Paisley, David M Blei, and Hanna Wallach. Bayesian poisson tensor factorization for inferring multilateral relations from sparse dyadic event counts. In \*SIGKDD\*, pp. 1045–1054, 2015.](#)
- [Mikkel N Schmidt, Ole Winther, and Lars Kai Hansen. Bayesian non-negative matrix factorization. In \*ICA 2009\*, pp. 540–547. Springer, 2009.](#)
- [Gordon K Smyth. Limma: linear models for microarray data. In \*Bioinformatics and computational biology solutions using R and Bioconductor\*, pp. 397–420. Springer, 2005.](#)
- [Harald Steck. Embarrassingly shallow autoencoders for sparse data. In \*WWW\*, pp. 3251–3257, 2019.](#)
- [David Michael Titterton, Adrian FM Smith, and Udi E Makov. \*Statistical analysis of finite mixture distributions\*. 1985.](#)
- [Jakub Tomczak and Max Welling. VAE with a VampPrior. In \*AISTATS\*, pp. 1214–1223. PMLR, 2018.](#)
- [Martin J Wainwright and Michael I Jordan. \*Graphical models, exponential families, and variational inference\*. Now Publishers, Inc., 2008.](#)
- [Wei Wang and Matthew Stephens. Empirical Bayes matrix factorization. \*JMLR\*, 22\(1\):5332–5371, 2021.](#)
- [Yixin Wang and David M. Blei. The Blessings of multiple causes. \*JASA\*, 114\(528\):1574–1596, 2019. Publisher: Taylor & Francis.](#)
- [Yixin Wang, David Blei, and John P Cunningham. Posterior collapse and latent variable non-identifiability. \*NeurIPS\*, 34:5443–5455, 2021.](#)
- [F Alexander Wolf, Philipp Angerer, and Fabian J Theis. SCANPY: large-scale single-cell gene expression data analysis. \*Genome biology\*, 19:1–5, 2018.](#)
- [Xinyi Zhong, Chang Su, and Zhou Fan. Empirical Bayes PCA in high dimensions. \*J. R. Stat. Soc.\*, 84\(3\): 853–878, 2022.](#)

## A Introduction

This is the supplement for the manuscript titled "Population priors for matrix factorization". The figures in this document supplement Figures 2, 3, and 4 in the main text. We provide derivations for twin population priors for Gaussian matrix factorization, as well as computing the likelihood of the held-out data in section ???. We then give more details about our experimental setup, including the parameters used in training (e.g., batch-size) in section ??. Finally, we provide results for additional experiments in section ???: for additional values of the latent dimension  $L$  for the datasets studied in the main text.

Note that this document is accompanied by an archive file `code.zip`, that contains the source code, instructions to install and run the code, and scripts to recreate the experiments and plot the figures in the manuscript. All scripts have been de-identified.

## B Derivations

In this section, we give more details and derivations for the quantities defined in the main text. Specifically, in section ?? we derive the twin population priors for the Gaussian matrix factorization (GMF). In section ??, we derive the variational approximation to the twin population priors in GMF. Finally, in section ??, we derive the expression for held-out likelihood in matrix factorization.

### B.1 Twin population priors for Gaussian matrix factorization

In this section we derive population priors for Gaussian matrix factorization (GMF). In the classical GMF, the likelihood and priors on the latent variables are Gaussian (?). The generative model is as follows:

$$\begin{aligned} W_{j,l} &\sim \mathcal{N}(0, \sigma_W^2), & j = 1 \dots D, l = 1 \dots L, \\ Z_{i,l} &\sim \mathcal{N}(0, \sigma_Z^2), & i = 1 \dots N, l = 1 \dots L, \\ X_{i,j} \mid \mathbf{U}_i, \mathbf{V}_j &\sim \mathcal{N}\left(\sum_{l=1}^L Z_{i,l} W_{j,l}, \sigma^2\right), & i = 1 \dots N, j = 1 \dots D, \end{aligned} \quad (28)$$

where  $\mathbf{U} = [Z_{i,l}] \in \mathbb{R}^{N \times L}$  and  $[\mathbf{W}_{j,l}] \in \mathbb{R}^{D \times L}$ , with  $\mathbf{U}_i$  and  $\mathbf{V}_j$  and are  $L$ -vectors representing the row and column wise latent variables, and  $\sigma$ ,  $\sigma_W$  and  $\sigma_Z$  are constant.

For TwinEB, our goal is to learn the prior for the row and column latent variables, so, as in the main text, we construct mixture priors (Equations ?? and ??), where  $P_{\boldsymbol{\mu}_k, \boldsymbol{\sigma}_k}$  and  $P_{\boldsymbol{\nu}_k, \boldsymbol{\eta}_k}$  are Gaussian parameterized by their mean and scale parameters. That is,

$$P_{\theta^r}(\mathbf{U}_i) = \sum_{k=1}^{K_r} \pi \omega_k \mathcal{N}(\mathbf{U}_i; \boldsymbol{\mu}_k, \boldsymbol{\sigma}_k), \quad (29)$$

$$P_{\theta^c}(\mathbf{V}_j) = \sum_{k=1}^{K_c} \rho_k \mathcal{N}(\mathbf{V}_j; \boldsymbol{\nu}_k, \boldsymbol{\eta}_k), \quad (30)$$

$$(31)$$

where  $K^r$  and  $K^c$  are the number of components in the mixtures,  $\theta^r = \{\boldsymbol{\mu}, \boldsymbol{\sigma}, \boldsymbol{\pi}\}$ ,  $\theta^c = \{\boldsymbol{\nu}, \boldsymbol{\eta}, \boldsymbol{\rho}\}$ . The locations  $\boldsymbol{\mu} \in \mathbb{R}^{K_r \times L}$  and  $\boldsymbol{\nu} \in \mathbb{R}^{K_c \times L}$ , the scales  $\boldsymbol{\sigma} \in \mathbb{R}^{K_r \times L}$  and  $\boldsymbol{\eta} \in \mathbb{R}^{K_c \times L}$ , and the mixture weights  $\boldsymbol{\pi} \in \Delta^{K_r}$ ,  $\boldsymbol{\omega} \in \Delta^{K_r}$  and  $\boldsymbol{\rho} \in \Delta^{K_c}$  are the parameters of the mixtures priors.

### B.2 Posterior inference in GMF with twin population priors

Our goal is to calculate the posterior  $P(\mathbf{U}, \mathbf{V} \mid \mathbf{X}, \theta^r, \theta^c)$ , given data  $\mathbf{X}$ , which also depends on our choice of priors  $\theta^r, \theta^c$ . Our strategy is the same as in the main text, namely, to simultaneously optimize for the parameters of the variational distribution and the TwinEB prior. We substitute the variational family in

Equations (??) and (??) with Gaussian, as follows:

$$\boldsymbol{\lambda}_{i,l}^r := (a'_{i,l}, b'_{i,l}), \quad (32)$$

$$\boldsymbol{\lambda}_{j,l}^c := (a_{j,l}, b_{j,l}), \quad (33)$$

$$q_{\boldsymbol{\lambda}_{i,l}^r}(U_{i,l}) := \mathcal{N}(a'_{i,l}, b'_{i,l}), \quad (34)$$

$$q_{\boldsymbol{\lambda}_{j,l}^c}(V_{j,l}) := \mathcal{N}(a_{j,l}, b_{j,l}). \quad (35)$$

Each Gaussian is parameterized by its natural parameters  $a$  and  $b$ :

$$\mathcal{N}(x; a, b) \propto \exp(ax + bx^2)$$

The log-likelihood of the data is:

$$\begin{aligned} \log P(\mathbf{X} \mid \mathbf{Z}, \mathbf{W}) &= \log \prod_{i=1}^N \prod_{j=1}^D P_{\sigma}(X_{i,j} \mid \mathbf{Z}_i, \mathbf{W}_j) \\ &= \sum_{i=1}^N \sum_{j=1}^D \log \mathcal{N}\left(X_{i,j}; \sum_{l=1}^L Z_{i,l} W_{j,l}, \sigma^2\right) \\ &= - \sum_{i=1}^N \sum_{j=1}^D \left[ \log \sigma \sqrt{2\pi} + \frac{1}{2} \left( \frac{X_{i,j} - \sum_{l=1}^L Z_{i,l} W_{j,l}}{\sigma} \right)^2 \right]. \end{aligned} \quad (36)$$

### B.3 Held-out likelihood for matrix factorization

We use the log likelihood of held-out data as the score for each model. Let  $\#\{X_{i,j}\}$  denote  $N^{\text{out}}$  entries of the held-out rows that were masked. We compute the likelihood of the masked entries via the posterior predictive distribution:

$$\log P(\not\! / X^{\text{in}}) = \log \prod_{i=1}^{N^{\text{out}}} P(i \mid X^{\text{in}}) \quad (37)$$

$$= \sum_{i=1}^{N^{\text{out}}} \log P(i \mid X^{\text{in}}). \quad (38)$$

For Bayesian matrix factorization we expand the summand in Equation (??) as follows:

$$\log P(i \mid X^{\text{in}}) \approx \log \iint P(i \mid \mathbf{U}, \mathbf{V}) P(\mathbf{U}, \mathbf{V} \mid X^{\text{in}}) d\mathbf{U} d\mathbf{V} \quad (39)$$

$$\approx \log \iint P(i \mid \mathbf{U}, \mathbf{V}) q_{\Lambda}(\mathbf{U}, \mathbf{V} \mid X^{\text{in}}) d\mathbf{U} d\mathbf{V} \quad (40)$$

$$= \log \mathbb{E}_{\mathbf{U}, \mathbf{V} \sim q_{\Lambda}(\cdot)} [P(i \mid \mathbf{U}, \mathbf{V})] \quad (41)$$

$$\approx \log \frac{1}{M} \sum_m^M P(i \mid \mathbf{U}^{(M)}, \mathbf{V}^{(M)}), \quad (42)$$

where  $M$  is the number of Monte Carlo samples from  $q_{\Lambda}(\mathbf{U}, \mathbf{V} \mid X^{\text{in}})$ . In Equation (??) we approximate the true posterior  $P(\mathbf{U}, \mathbf{V} \mid X^{\text{in}})$  with its variational counterpart  $q_{\Lambda}$ . The log-likelihood score for the entire held-out data is then:

$$\sum_i^{N^{\text{out}}} \log P(i \mid X^{\text{in}}) \approx \sum_i^{N^{\text{out}}} \log \frac{1}{M} \sum_m^M P(i \mid \mathbf{U}^{(M)}, \mathbf{V}^{(M)}). \quad (43)$$

## C Experimental details

In this section we give more details on our experimental studies in the main manuscript. In section ?? we describe preprocessing for the gene expression in the Ru1322-scRNAseq dataset. In section ??, we specify the parameters used during training. In section ?? we give a brief description of the artifacts that accompany this supplementary material.

### C.1 Preprocessing of the Ru1322-scRNAseq gene expression dataset

We use CellRanger version 6.0.1 to process the FASTQ and generate the unique molecular identifier (UMI) count matrices. We use `scanpy` to preprocess the data, and the `seuratv3` algorithm to select highly variable genes ?.

### C.2 Training details

We set a batch size of 128 in all our experiments. We ran Poisson and Gaussian matrix factorization experiments for a maximum of 20,000 iterations. By this step, all runs had converged.

We initialized the learning rate for the row and column variables, `r1r` and `c1r` separately. We fix the initial learning rate `r1r`  $\in$  {0.01} and `c1r`  $\in$  {0.01}. In the experiments in the main text, we use 10 Monte Carlo samples to approximate the ELBO, while in the supplemental experiments, we use a single particle.

For the PMF experiments in the supplement, we subsample zeros as is standard (?). We uniformly randomly subsample the same number of zeros as non-zero values to estimate the likelihood. Concretely, let  $\mathcal{L}()$  denote the log-likelihood of the masked entries of the held-out rows,  $\#\{X_{i,j}\}$ . Then  $\mathcal{L}()$  can be decomposed as the sum of non-zero and zero  $X_{i,j}$ :

$$\mathcal{L}(\gamma) = \mathcal{L}(x_{i,j} \neq 0) + \mathcal{L}(x_{i,j} = 0), \quad (44)$$

where  $x_{i,j} \neq 0$  and  $x_{i,j} = 0$  have cardinalities  $N_{\text{non-zero}}^{\text{out}}$  and  $N_{\text{zero}}^{\text{out}}$  respectively.

We approximate Equation (??) by subsampling  $N_{\text{sub}} = \min(N_{\text{zero}}^{\text{out}}, N_{\text{non-zero}}^{\text{out}})$  of the zeros. Let  $\mathbf{X}_{N_{\text{sub}}}^{\text{zero}}$  denote a multiset of zeros of cardinality  $N_{\text{sub}}$ , then:

$$\hat{\mathcal{L}}(\gamma) \approx \mathcal{L}(x_{i,j} \neq 0) + \frac{N_{\text{zero}}^{\text{out}}}{N_{\text{sub}}} \mathbf{X}_{N_{\text{sub}}}^{\text{zero}}. \quad (45)$$

We ran our experiment on a machine equipped with an NVIDIA A100 GPU with 80GB memory. We implemented all methods in pytorch (?).

### C.3 Implementation

Please refer to the `code` directory for the source code and instructions on how to run the model. The file `README.md` contains instructions for installing the software, and running it. Under the `data` directory, we included preprocessed data for the MovieLens-100K dataset, a smaller version of the MovieLens-1M studied in the main text. In the `notebooks` directory, we put notebooks used for preprocessing the data, and plotting the figures in the manuscript. Finally, in the `pipelines` directory, we put `nextflow` scripts that recreate experiments that we have run for the manuscript.

## D Additional experiments

In this section we present additional experimental results. We show results for additional values of the latent dimension  $L$ . Concretely, section ?? studies the effects of the twin population priors on Poisson and Gaussian matrix factorization, on three real world datasets, while section ?? examines the sensitivity of the

twin population priors to the choice of its hyper-parameters. We find that these results corroborate those that were presented in the main manuscript, that is, matrix factorization with traditional priors is sensitive to the choice of the hyper-parameters of the prior, and twin population priors is a robust way to set the prior in this family of models.

### D.1 Additional values of latent dimensions

In the main text, we study the effects of twin population priors on the Poisson and Gaussian matrix factorization over four real datasets, with the dimension of the latent variables set to  $L = 15$ . Here we present results for additional values of  $L$ . Figures ??, ??, ??, ??, and ?? show the results. However, the twin population priors find a comparable or better prior, in the sense of higher held-out likelihood.

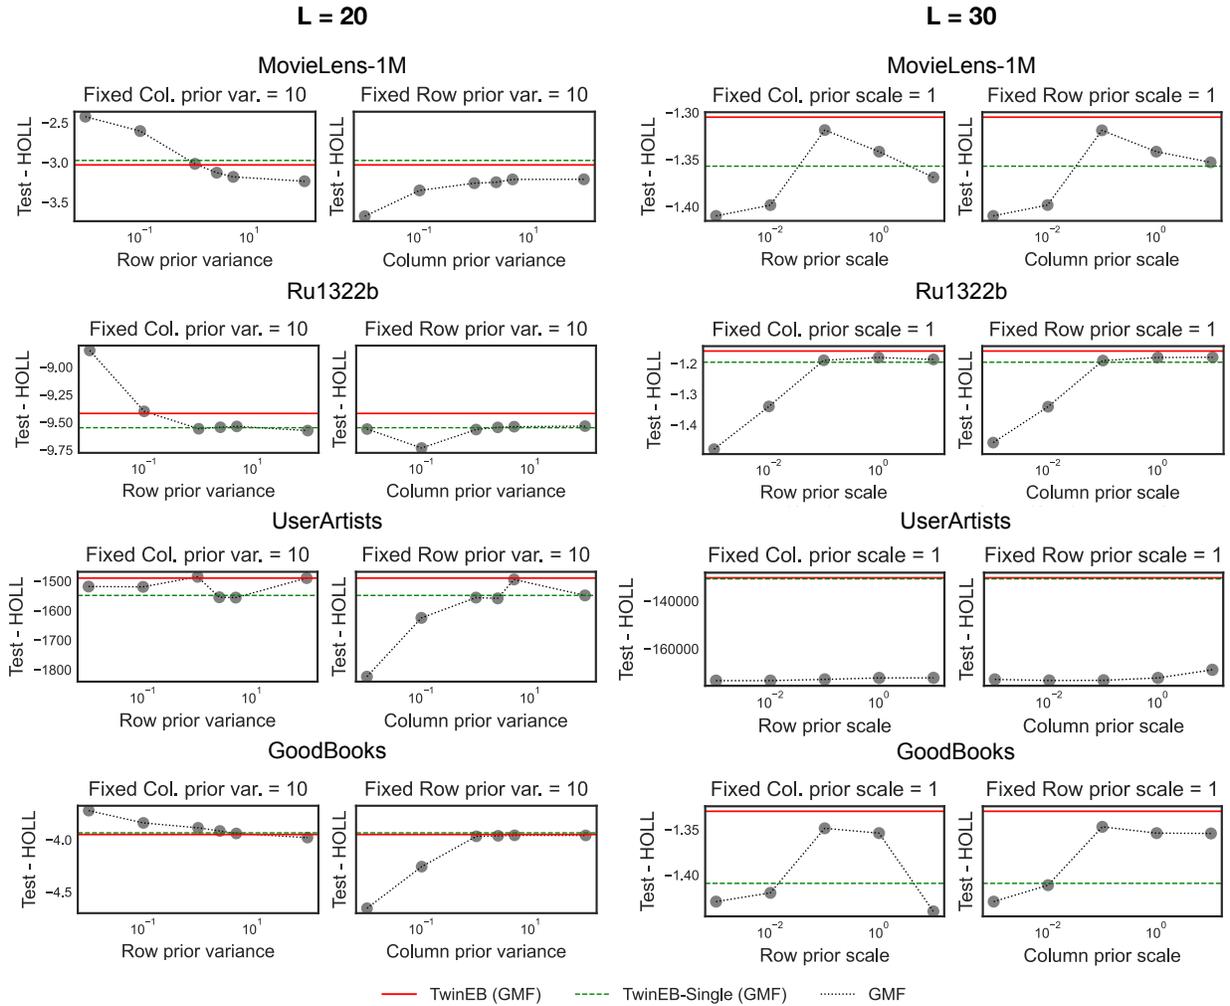


Figure 5: As in Figure ??, but with additional values of latent dimension  $L$ .

### D.2 Performance Increase with Number of Mixture Components

We show the performance of TwinEB on simulated data as we vary the number of row and column mixture components. Here we add results for additional values of  $L$ , and also study the MovieLens 100K dataset. The result in Figures ?? suggests that as the number of mixture components increases, the performance of TwinEB improves. The gain is apparent when the number of row components  $K_r$  increases. This is expected, as our evaluation procedure, measures the generalizability for held-out rows, but not held-out columns.

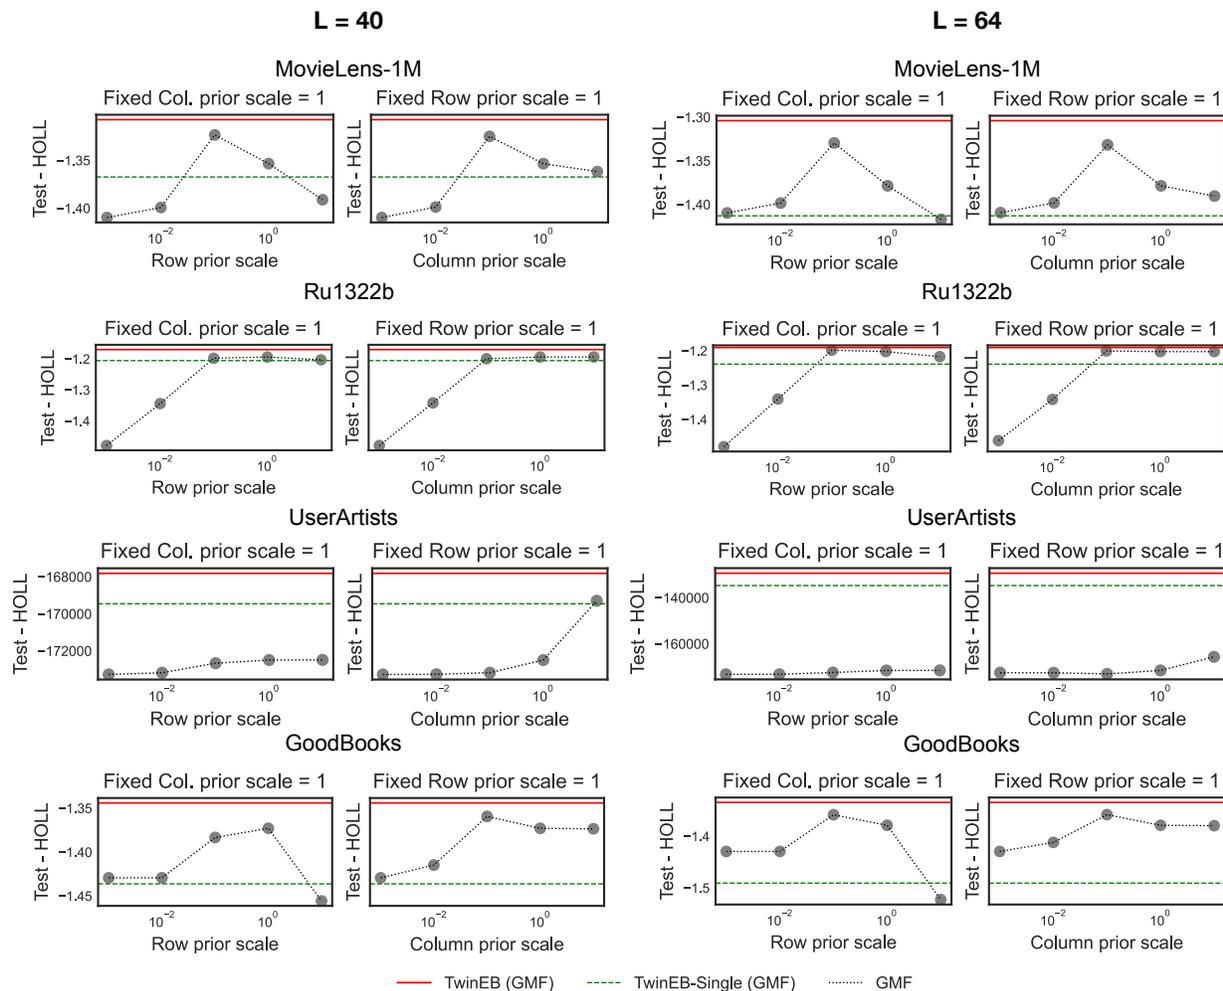
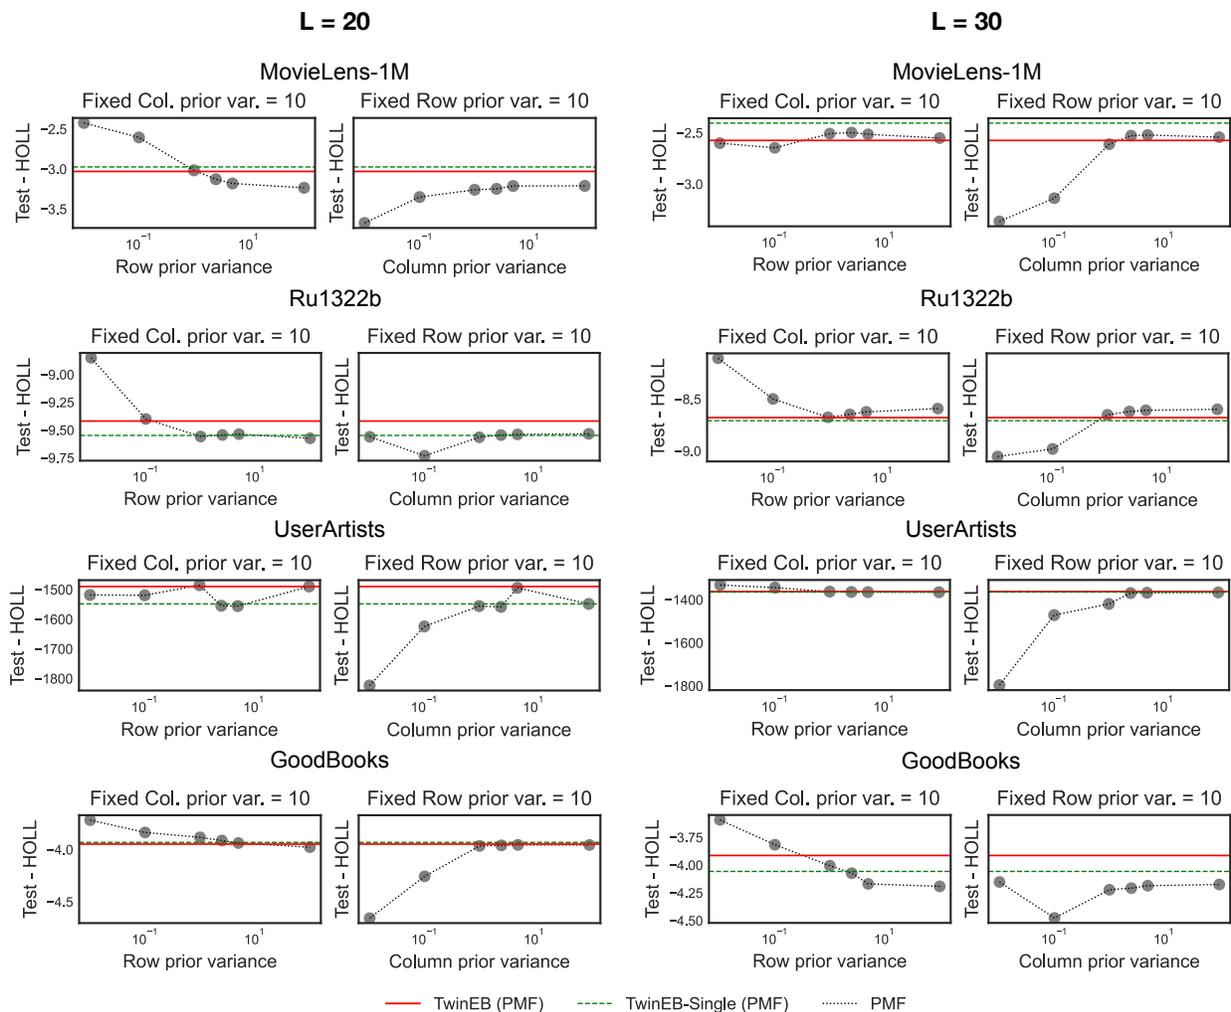
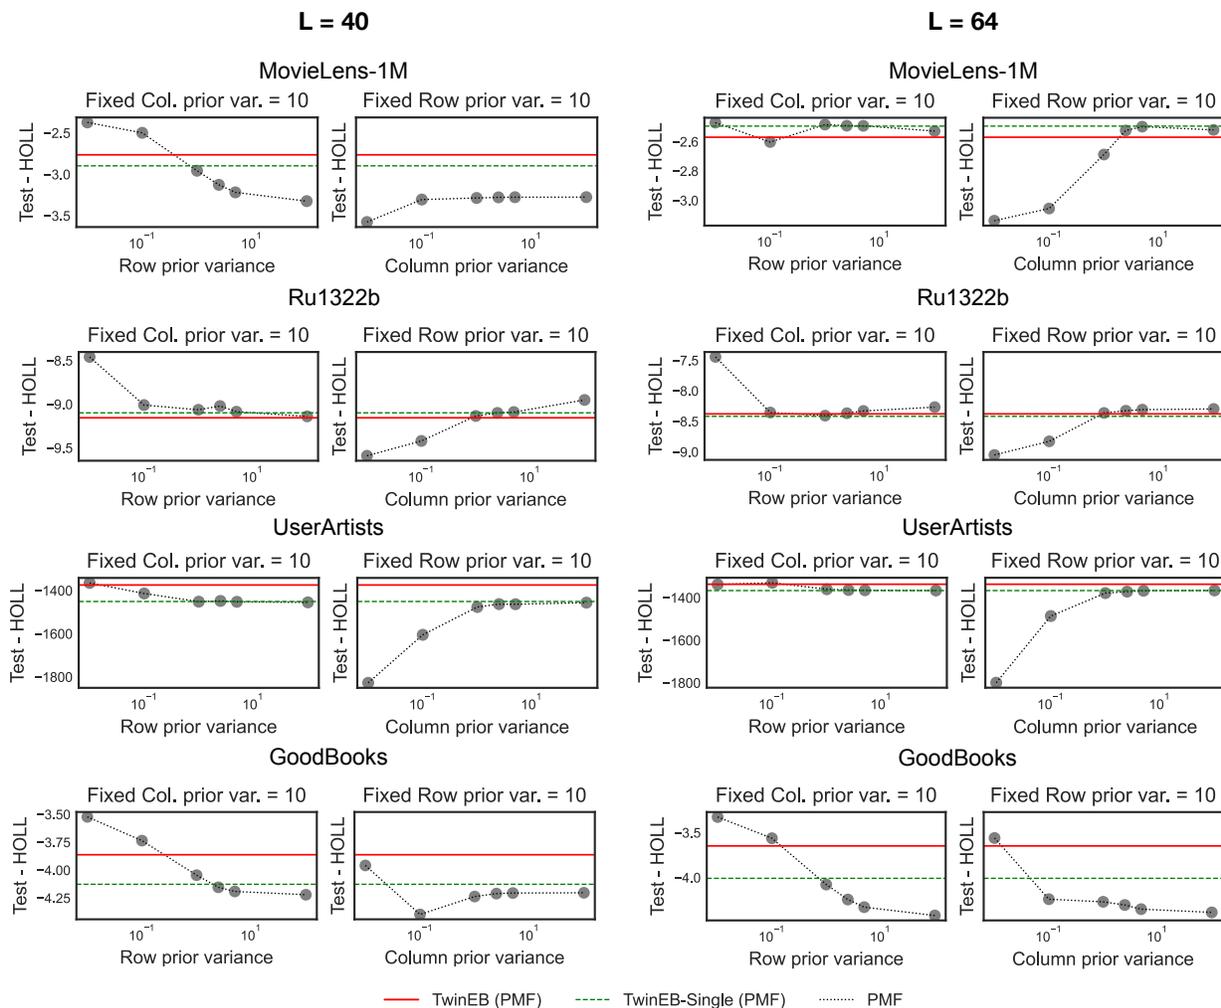


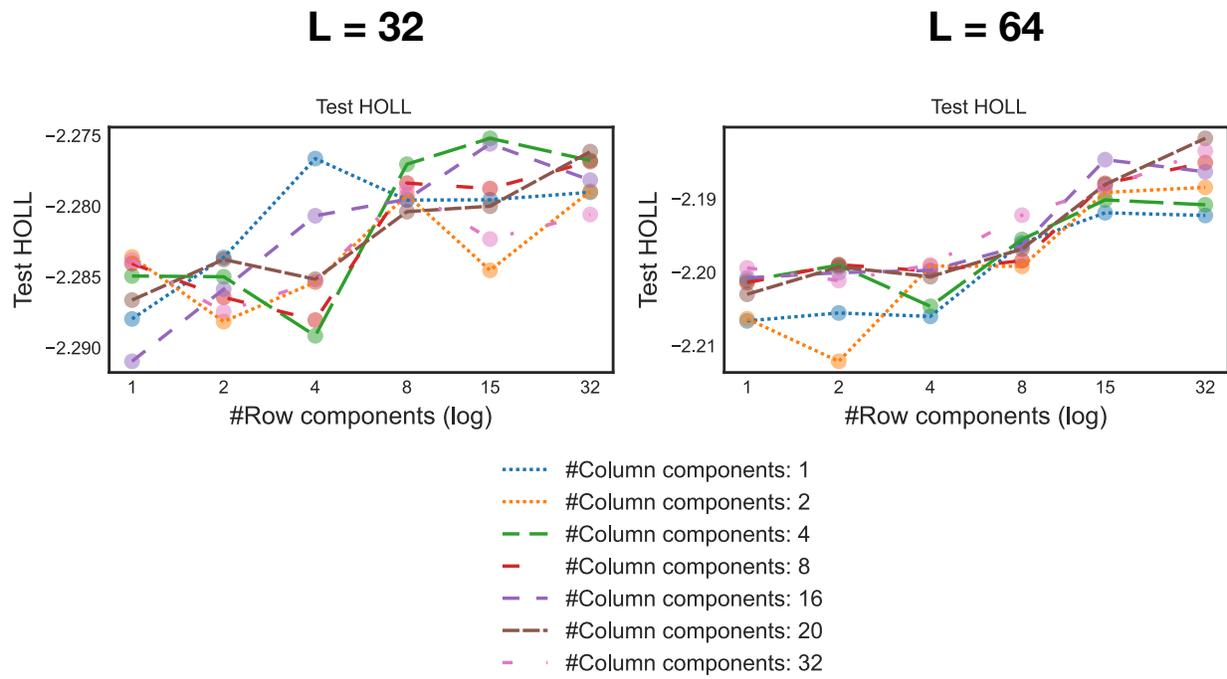
Figure 6: As in Figure ??, but with additional values of latent dimension  $L$ .



**Figure 7:** As in Figure ??, but with additional values of latent dimension  $L$ .



**Figure 8:** As in Figure ??, but with additional values of latent dimension  $L$ .



**Figure 9:** As in Figure ??, but with additional column components as well as latent dimension  $L = 32$ .