
This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Qin, Xiangju; Blomstedt, Paul; Leppäaho, Eemeli; Parviainen, Pekka; Kaski, Samuel
Distributed Bayesian matrix factorization with limited communication

Published in:
Machine Learning

DOI:
[10.1007/s10994-019-05778-2](https://doi.org/10.1007/s10994-019-05778-2)

Published: 01/10/2019

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Qin, X., Blomstedt, P., Leppäaho, E., Parviainen, P., & Kaski, S. (2019). Distributed Bayesian matrix factorization with limited communication. *Machine Learning*, 108(10), 1805–1830.
<https://doi.org/10.1007/s10994-019-05778-2>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.



Distributed Bayesian matrix factorization with limited communication

Xiangju Qin¹ · Paul Blomstedt¹ · Eemeli Leppäaho¹ · Pekka Parviainen^{1,2} · Samuel Kaski¹

Received: 28 January 2018 / Accepted: 8 January 2019 / Published online: 25 March 2019

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2019

Abstract

Bayesian matrix factorization (BMF) is a powerful tool for producing low-rank representations of matrices and for predicting missing values and providing confidence intervals. Scaling up the posterior inference for massive-scale matrices is challenging and requires distributing both data and computation over many workers, making communication the main computational bottleneck. Embarrassingly parallel inference would remove the communication needed, by using completely independent computations on different data subsets, but it suffers from the inherent unidentifiability of BMF solutions. We introduce a hierarchical decomposition of the joint posterior distribution, which couples the subset inferences, allowing for embarrassingly parallel computations in a sequence of at most three stages. Using an efficient approximate implementation, we show improvements empirically on both real and simulated data. Our distributed approach is able to achieve a speed-up of almost an order of magnitude over the full posterior, with a negligible effect on predictive accuracy. Our method outperforms state-of-the-art embarrassingly parallel MCMC methods in accuracy, and achieves results competitive to other available distributed and parallel implementations of BMF.

Xiangju Qin and Paul Blomstedt have contributed equally to this work.

Editors: Jesse Davis, Elisa Fromont, Derek Greene, and Bjorn Bringmann.

✉ Xiangju Qin
xiangju.qin@aalto.fi

Paul Blomstedt
paul.blomstedt@aalto.fi

Eemeli Leppäaho
eemeli.leppaaho@aalto.fi

Pekka Parviainen
Pekka.Parviainen@uib.no

Samuel Kaski
samuel.kaski@aalto.fi

¹ Department of Computer Science, Helsinki Institute for Information Technology (HIIT), Aalto University, Konemiehentie 2, Espoo, Finland

² Present Address: Department of Informatics, University of Bergen, Bergen, Norway

Keywords Bayesian matrix factorization · Embarrassingly parallel MCMC · Distributed inference · Posterior propagation

1 Introduction

Latent variable models based on matrix factorization have in recent years become one of the most popular and successful approaches for matrix completion tasks, such as collaborative filtering in recommender systems (Koren et al. 2009) and drug discovery (Cobanoglu et al. 2013). The main idea in such models is, given a matrix of observed values $\mathbf{Y} \in \mathbb{R}^{N \times D}$, to find two matrices $\mathbf{X} \in \mathbb{R}^{N \times K}$ and $\mathbf{W} \in \mathbb{R}^{D \times K}$ with $K \ll N, D$, such that their product forms a low-rank approximation of \mathbf{Y} :

$$\mathbf{Y} \approx \mathbf{X}\mathbf{W}^T. \quad (1)$$

In matrix completion, the matrix \mathbf{Y} is typically very sparsely observed, and the goal is to predict unobserved matrix elements based on the observed ones.

A standard way of dealing with high levels of unobserved elements in matrix factorization is to model the observed values only (instead of imputing the missing values), using regularization to avoid overfitting. A probabilistically justified regularized matrix factorization model was first introduced by Salakhutdinov and Mnih (2008b), and subsequently extended to a fully Bayesian formulation (called Bayesian probabilistic matrix factorization, BPMF) (Salakhutdinov and Mnih 2008a). This formulation sidesteps the difficulty of choosing appropriate values for the regularization parameters by considering them as hyperparameters, placing a hyperprior over them and using Markov chain Monte Carlo (MCMC) to perform posterior inference. Additional advantages of the fully Bayesian approach include improved predictive accuracy, quantification of the uncertainty in predictions, the ability to incorporate prior knowledge, as well as flexible utilization of side-information (Adams et al. 2010; Park et al. 2013; Porteous et al. 2010; Simm et al. 2015).

Given the appeal and many advantages of Bayesian matrix factorization, applying it also to massive-scale matrices would be attractive but scaling up the posterior inference has proven difficult, and calls for distributing both data and computation over many workers. So far only very few distributed implementations of BMF have been presented in the literature. Recently, Ahn et al. (2015) proposed a solution based on distributed stochastic gradient Langevin dynamics (DSGLD), and showed empirically that BMF with DSGLD achieves the same level of predictive performance as Gibbs sampling. However, the convergence efficiency of the DSGLD solution is constrained by several factors such as the need for careful tuning of the learning rate ϵ_t and for using an orthogonal group partition¹ for training. When a model is trained with blocks in an orthogonal group, in each iteration it only makes use of a small subset of the full data set for learning, which could lead to an estimate with higher variance and slowing down the convergence speed. Şimşekli et al. (2015) developed a similar distributed MCMC method based on SGLD for large generalised matrix factorization problems, which they called Parallel SGLD (PSGLD); see also Şimşekli et al. (2017) for an application of the method to non-negative matrix factorization (NMF). Different from DSGLD, PSGLD is implemented such that for each iteration, only blocks of \mathbf{W} instead of the whole \mathbf{W} need to be transferred among parallel workers. Nevertheless, this solution suffers from the same issues as DSGLD. Vander Aa et al. (2017) presented a distributed

¹ This partition is used to avoid conflicting access to parameters among parallel workers, and refers to a partition scheme in which rows and columns included in one block do not appear in the other ones.

high-performance implementation of BPMF with Gibbs sampling using the TBB and GASPI libraries, and provided an empirical comparison with other state-of-the-art distributed high-performance parallel implementations. They found that a significant speed-up could only be achieved with a limited number of workers, after which the addition of more workers eventually leads to a dramatic drop in parallel computation efficiency due to the increased communication overhead (Vander Aa et al. 2016, 2017). Therefore, a key factor in devising even more scalable distributed solutions is to be able to minimize communication between worker nodes.

One of the most promising directions in large-scale Bayesian computation in recent years has been *embarrassingly parallel* MCMC, a family of essentially communication-free algorithms, where the data are first partitioned into multiple subsets and independent sampling algorithms are then run on each subset in parallel (Minsker et al. 2014; Neiswanger et al. 2014; Scott et al. 2016; Srivastava et al. 2015; Wang and Dunson 2013; Wang et al. 2014, 2015). In these algorithms, communication only takes place in a final aggregation step, which combines the subset posteriors to form an approximation to the full-data posterior. A key factor which limits the applicability of such methods for BMF models is that Eq. (1) can be solved only up to orthogonal transformations. Each subset posterior can therefore converge to any of an infinite number of modes, making the aggregation step difficult to carry out in a meaningful way. Previous embarrassingly parallel MCMC algorithms have only been applied in cases where the model is unidentified up to a finite number of solutions (e.g. Nemeth and Sherlock 2017) but are not applicable in a continuum of unidentifiable cases.

In this paper, we introduce an approach which addresses the unidentifiability issue by introducing dependencies between the subset posterior inferences, while limiting the communication between workers. We will draw inspiration from the observation that even though an infinite number of solutions to Eq. (1) exist in principle, in practical computation with a finite number of observations, a sampler with finite chain-length will only explore a small number of solutions, each corresponding to a separate mode. The key idea is to encourage the samplers in all subsets to target the same set of solutions; note that this does not restrict generality as the standard way of finding other modes, by employing additional chains, is available here as well.

In large BMF problems we partition the data matrix \mathbf{Y} along both rows and columns, effectively making different subsets dependent on parameters shared by subsets on the same rows and columns. To implement the dependencies in the inference, we divide the sub-

Fig. 1 An example illustrating posterior propagation (PP) for a data matrix \mathbf{Y} partitioned into 3×4 subsets. Subset inferences proceed in three successive stages, with posteriors obtained in one stage being propagated as priors to the next one; the numbers in the \mathbf{Y} matrix denote the stage in which the particular subset is processed. Within each stage, the subsets are processed in parallel with no communication

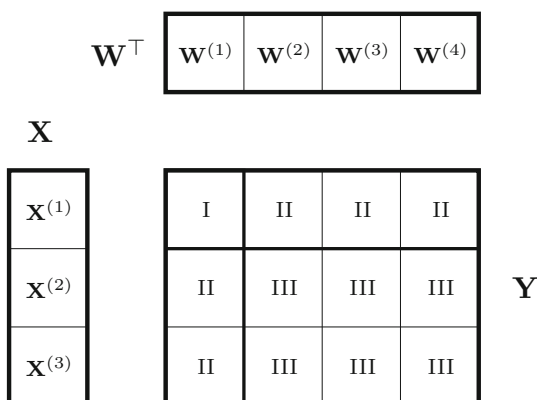


Table 1 Comparison of different implementations of parallel and distributed BMF/NMF for a data matrix $\mathbf{Y} \in \mathbb{R}^{N \times D}$

Models	Tune learning rate or hyperparam.	Shuffle train data per iter.	Optimal partition	Communication cost
BMF + DSGLD (Ahn et al. 2015)	✓	✓	Squared orthogonal	$\mathcal{O}((N + D)K(\sqrt{U} + 1)T)$
NMF + PSGLD (Şimşekli et al. 2017)	✓	✓	Squared orthogonal	$\mathcal{O}(DKT)$
Distributed BPFM (Vander Aa et al. 2017)	–	–	Load balance	$\mathcal{O}((N + D)K(U - 1)T)$
Proposed method	–	–	Flexible	$\mathcal{O}((N + D)(K + K^2)\sqrt{U})$

We assume a maximum of U compute nodes available for distributed inference and a partition scheme of $(\sqrt{U} + 1) \times (\sqrt{U} + 1)$ for all methods. The number of latent dimensions and iterations are K and T , respectively. For simplicity, \mathbf{Y} is assumed to be fully observed. For our method, the K^2 term is due to using a full covariance matrix for the posterior of parameters \mathbf{W} and \mathbf{X} . For BMF + DSGLD, the $\sqrt{U} + 1$ term is applied only when $\sqrt{U} + 1$ parallel chains are deployed. Note that communication cost is incurred only when the model is running on multiple compute nodes (distributed memory). Also note that the cost of loading the data is ignored; there is no communication incurred during inference for the embarrassingly parallel methods

sets into three groups, which are processed in a hierarchy of three consecutive stages (see Fig. 1). The posterior distributions obtained in each stage are propagated forwards and used as priors in the following stage. This way, communication only takes place between the stages and not between the subsets within a stage, and each subset inference is regularized using information from the relevant subset inference in the preceding stage. Note that within each stage, we perform the inference for subsets in parallel. Thus, for a partition scheme with $r \times c$ subsets, the maximum number of parallel workers that can be used by our algorithm is equal to the number of subsets in the third stage, i.e. $(r - 1) \times (c - 1)$. We refer to the proposed procedure as *posterior propagation (PP)*. Table 1 compares the computational characteristics of PP with previous approaches for parallel and distributed BMF/NMF.

1.1 Contributions and overview of the paper

The main contributions of our paper are as follows: In Sect. 3, we introduce a hierarchical, exact decomposition of the joint posterior of the BMF model parameters, which makes possible embarrassingly parallel computations over data subsets in a sequence of at most three stages, limiting all communication to take place between the stages. This decomposition is computationally intractable in general; however, in Sect. 4 we build on it to develop a MCMC-based approximate inference scheme for BMF. In the numerical experiments of Sect. 5, we show empirically, with both real and simulated data, that the proposed distributed approach is able to achieve a speed-up of almost an order of magnitude over the full posterior, with a negligible effect on predictive accuracy, compared to MCMC inference on the full data. In the experiments, the method also significantly outperforms state-of-the-art embarrassingly parallel MCMC methods in accuracy, and achieves competitive results compared to other available distributed and parallel implementations of BMF.

2 Background

2.1 Bayesian matrix factorization

Let $\mathbf{Y} \in \mathbb{R}^{N \times D}$ be a partially observed data matrix, $\mathbf{X} \in \mathbb{R}^{N \times K} = (\mathbf{x}_1, \dots, \mathbf{x}_N)^\top$ and $\mathbf{W} = (\mathbf{w}_1, \dots, \mathbf{w}_D)^\top \in \mathbb{R}^{D \times K}$ be matrices of unknown parameters. The general Bayesian matrix factorization (BMF) model is then specified by the likelihood

$$p(\mathbf{Y}|\mathbf{X}, \mathbf{W}) = \prod_{n=1}^N \prod_{d=1}^D \left[p(y_{nd} | \mathbf{x}_n^\top \mathbf{w}_d) \right]^{\mathbf{1}_{nd}}, \quad (2)$$

which is a probabilistic version of Eq. (1). Here $\mathbf{1}_{nd}$ denotes an indicator function which equals 1 if the element y_{nd} is observed and 0 otherwise.

While the general BMF model is agnostic to the choice of distributional form, in many applications, the elements y_{nd} of the data matrix are assumed to be normally distributed, conditionally on the parameter vectors \mathbf{x}_n and \mathbf{w}_d ,

$$p(y_{nd} | \mathbf{x}_n^\top \mathbf{w}_d) = \mathcal{N}(y_{nd} | \mathbf{x}_n^\top \mathbf{w}_d, \tau^{-1}), \quad (3)$$

where τ denotes the noise precision. Note that some formulations specify an individual precision τ_d for each column. To complete the Bayesian model, priors are placed on the model parameters \mathbf{X} and \mathbf{W} , commonly normal priors specified as

$$p(\mathbf{X} | \mu_{\mathbf{X}}, \Lambda_{\mathbf{X}}^{-1}) = \prod_{n=1}^N \mathcal{N}_K(\mathbf{x}_n | \mu_{\mathbf{X}}, \Lambda_{\mathbf{X}}^{-1}), \quad (4a)$$

$$p(\mathbf{W} | \mu_{\mathbf{W}}, \Lambda_{\mathbf{W}}^{-1}) = \prod_{d=1}^D \mathcal{N}_K(\mathbf{w}_d | \mu_{\mathbf{W}}, \Lambda_{\mathbf{W}}^{-1}), \quad (4b)$$

where \mathcal{N}_K denotes a K -dimensional normal distribution with covariance specified in terms of the precision matrix. The model formulation may additionally include priors on some or all of the hyperparameters $\mu_{\mathbf{X}}, \Lambda_{\mathbf{X}}, \mu_{\mathbf{W}}, \Lambda_{\mathbf{W}}$, as well as on the data precision parameters τ_d (e.g. Bhattacharya and Dunson 2011; Salakhutdinov and Mnih 2008a). For concreteness, we proceed in this paper with the Gaussian case, as specified in Eqs. (3–4b), but note that the developments of Sect. 3, along with Algorithm 1, are general with no reference to this choice of distributions for likelihood and priors.

2.1.1 Unidentifiability of matrix factorization

It is commonly known that the solution of Eq. (1) is unique only up to orthogonal transformations. To demonstrate this unidentifiability, let \mathbf{P} be any semi-orthogonal matrix for which $\mathbf{P}\mathbf{P}^\top = \mathbf{I}_K$, \mathbf{I}_K being a $K \times K$ unit matrix, and denote $\hat{\mathbf{Y}} := \mathbf{X}\mathbf{W}^\top$. Then, performing an orthogonal transformation by right-multiplying both \mathbf{X} and \mathbf{W} by \mathbf{P} , leads to

$$\mathbf{X}\mathbf{P}(\mathbf{W}\mathbf{P})^\top = \mathbf{X}\mathbf{P}\mathbf{P}^\top\mathbf{W}^\top = \mathbf{X}\mathbf{W}^\top = \hat{\mathbf{Y}},$$

by which an uncountable number of equally good solutions to Eq. (1) can be produced.

As a special case, let \mathbf{P} be a $K \times K$ unit matrix with the k th diagonal element set to -1 . The matrix \mathbf{P} is then clearly orthogonal (also semi-orthogonal), since $\mathbf{P}^\top\mathbf{P} = \mathbf{P}\mathbf{P}^\top = \mathbf{I}_K$. Right-multiplying \mathbf{X} and \mathbf{W} by \mathbf{P} has the effect of flipping the signs of all elements in the

k th columns of these matrices. It can then easily be verified that the product of the resulting matrices remains unchanged. Since any of the K columns of \mathbf{X} and \mathbf{W} can have their signs flipped without affecting the product $\hat{\mathbf{Y}}$, within this family of transformations we have 2^K equally good solutions for Eq. (1).

Although the unidentifiability of a single matrix factorization task could be addressed by e.g. constraining \mathbf{W} to be a lower triangular matrix (Lopes and West 2004), expensive communication would be needed for distributed inference schemes to ensure all the posteriors are jointly identifiable.

2.2 Embarrassingly parallel MCMC

Consider now a parametric model $p(\mathbf{Y}|\theta)$ with exchangeable observations $\mathbf{Y} = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ and parameter θ for which we wish to perform posterior inference using MCMC. If N is very large, the inference may be computationally too expensive to be carried out on a single machine. Embarrassingly parallel MCMC strategies aim to overcome this by partitioning the data \mathbf{Y} into multiple disjoint subsets $\mathbf{Y}^{(1)} \cup \dots \cup \mathbf{Y}^{(J)} = \mathbf{Y}$, and running independent sampling algorithms for each subset using a down-weighted prior $p(\theta)^{1/J}$. In most embarrassingly parallel MCMC algorithms, the aggregation of the obtained subset posteriors into a full-data posterior is based, in one way or another, on the following *product density equation*:

$$p(\theta|\mathbf{Y}) \propto p(\theta)p(\mathbf{Y}|\theta) = \prod_{j=1}^J p(\theta)^{1/J} p(\mathbf{Y}^{(j)}|\theta), \quad (5)$$

where each factor $p(\theta)^{1/J} p(\mathbf{Y}^{(j)}|\theta)$ constitutes an independent inference task. Aggregating the joint *product density equation* with satisfactory efficiency and accuracy is in general a challenging problem, since the involved densities are unknown and represented in terms of samples. For a recent overview of various subset posterior aggregation techniques, see Angelino et al. (2016).

Standard embarrassingly parallel inference techniques relying on Eq. (5) are not well-suited for unidentifiable models, such as the BMF model presented in Sect. 2.1. To illustrate this, consider the following simple example:

Example 1 Assume that we have observed a data matrix

$$\mathbf{Y} = \begin{bmatrix} 1 & 4 & 16 \end{bmatrix},$$

conditional on which we wish to estimate the parameters \mathbf{X} and \mathbf{W} of the corresponding BMF model. A plausible inference may then result in a bimodal posterior with high density regions around, say, the exact solutions

$$\mathbf{X} = \begin{bmatrix} 4 \end{bmatrix}, \quad \mathbf{W} = \begin{bmatrix} 0.25 & 1 & 4 \end{bmatrix},$$

and $-\mathbf{X}$, $-\mathbf{W}$. Next, assume that we split the data into three subsets

$$\mathbf{Y}^{(1)} = \begin{bmatrix} 1 \end{bmatrix}, \quad \mathbf{Y}^{(2)} = \begin{bmatrix} 4 \end{bmatrix}, \quad \mathbf{Y}^{(3)} = \begin{bmatrix} 16 \end{bmatrix},$$

and conduct independent inference over each of them. Again, plausible subset inferences may accumulate posterior mass around *some* set of exact solutions, say,

$$\begin{aligned}\mathbf{X} &= [1], \quad \mathbf{W}^{(1)} = [1], \\ \mathbf{X} &= [2], \quad \mathbf{W}^{(2)} = [2], \\ \mathbf{X} &= [4], \quad \mathbf{W}^{(3)} = [4],\end{aligned}$$

along with their corresponding negative solutions. However, aggregating these inferences using Eq. (5) does not necessarily lead to a posterior with high density around any correct solution.

Ideally, we would like all subset inferences in the above example to target the same solutions in order for them to reinforce each other. To do so, it is clearly necessary to impose some constraints or regularization on them. One way of doing this is to equip the inferences with strong enough prior information. We will build on this idea in the following section.

3 Hierarchical parallelization of BMF

Let us now assume that a data matrix \mathbf{Y} has been partitioned with respect to both rows and columns into $I \times J$ subsets $\mathbf{Y}^{(i,j)}$, $i = 1, \dots, I$, $j = 1, \dots, J$. It then follows from Eqs. (2, 4a, 4b), that the joint posterior density of the BMF parameter matrices \mathbf{X} and \mathbf{W} , given the partitioned data matrix \mathbf{Y} , can be factorized as

$$\begin{aligned}p(\mathbf{X}, \mathbf{W}|\mathbf{Y}) &\propto p(\mathbf{X}) p(\mathbf{W}) p(\mathbf{Y}|\mathbf{X}, \mathbf{W}) \\ &= \prod_{i=1}^I p(\mathbf{X}^{(i)}) \prod_{j=1}^J p(\mathbf{W}^{(j)}) \prod_{i=1}^I \prod_{j=1}^J p(\mathbf{Y}^{(i,j)}|\mathbf{X}^{(i)}, \mathbf{W}^{(j)}).\end{aligned}\quad (6)$$

Our goal is to develop an equivalent decomposition that fulfils the apparently contradictory aims of both allowing for embarrassingly parallel computations and making the subset inferences dependent.

3.1 From sequential to parallel inference

We begin with the simple case of having only three subsets. With $I = 1$ and $J = 3$, the parameters of the partitioned BMF model are \mathbf{X} , $\mathbf{W}^{(1)}$, $\mathbf{W}^{(2)}$ and $\mathbf{W}^{(3)}$. Sequential inference (exploiting no parallelism) over \mathbf{Y} can then be performed in three successive stages as follows. In the first stage, the posteriors for the parameters \mathbf{X} and $\mathbf{W}^{(1)}$, given $\mathbf{Y}^{(1)}$, are computed as

$$p(\mathbf{X}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1)}) \propto p(\mathbf{X}) p(\mathbf{W}^{(1)}) p(\mathbf{Y}^{(1)}|\mathbf{X}, \mathbf{W}^{(1)}). \quad (7)$$

In the second stage, the posterior from the first stage is used as a prior for the shared parameter \mathbf{X} to compute

$$\begin{aligned}p(\mathbf{X}, \mathbf{W}^{(1)}, \mathbf{W}^{(2)}|\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}) &\propto p(\mathbf{X}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1)}) p(\mathbf{W}^{(2)}) \\ &\quad \times p(\mathbf{Y}^{(2)}|\mathbf{X}, \mathbf{W}^{(2)}).\end{aligned}\quad (8)$$

In the above stage, using the posterior obtained in Eq. (7) as a prior can be interpreted as a form of regularization, which encourages the inference to target the same set of modes as the first stage. Finally, using the posterior from the second stage as a prior in the third stage then gives the full-data posterior as

$$\begin{aligned}
 p(\mathbf{X}, \mathbf{W}|\mathbf{Y}) &= p(\mathbf{X}, \mathbf{W}^{(1)}, \mathbf{W}^{(2)}, \mathbf{W}^{(3)}|\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}, \mathbf{Y}^{(3)}) \\
 &\propto p(\mathbf{X}, \mathbf{W}^{(1)}, \mathbf{W}^{(2)}|\mathbf{Y}^{(1)}, \mathbf{Y}^{(2)}) p(\mathbf{W}^{(3)}) \\
 &\quad \times p(\mathbf{Y}^{(3)}|\mathbf{X}, \mathbf{W}^{(3)}).
 \end{aligned}$$

In general, a data set partitioned into J subsets will require J stages of sequential inference to obtain the full posterior.

We will now consider an alternative, partly parallelizable inference scheme, which begins with an initial stage identical to that of the above sequential scheme. However, instead of processing the subsets $\mathbf{Y}^{(2)}$ and $\mathbf{Y}^{(3)}$ in sequence, we process them in parallel. Regularizing the inferences with a common informative prior, obtained in the first stage, introduces dependence between them and encourages the targeted solutions to agree with each other. This leads to the following decomposition:

$$p(\mathbf{X}, \mathbf{W}|\mathbf{Y}) \propto p(\mathbf{X}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1)}) \quad (9a)$$

$$\times \left[p(\mathbf{X}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1)}) p(\mathbf{W}^{(2)}) p(\mathbf{Y}^{(2)}|\mathbf{X}, \mathbf{W}^{(2)}) \right] \quad (9b)$$

$$\times \left[p(\mathbf{X}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1)}) p(\mathbf{W}^{(3)}) p(\mathbf{Y}^{(3)}|\mathbf{X}, \mathbf{W}^{(3)}) \right] \quad (9c)$$

$$\times p(\mathbf{X}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1)})^{-2}, \quad (9d)$$

where the right-hand side of line (9a) corresponds to the first stage, lines (9b)–(9c) correspond to the second stage, with the two remaining subsets now being processed in parallel. Finally, an aggregation stage combines all of (9a)–(9d).

With $J = 2$, the number of stages for the parallel scheme is exactly the same as for the sequential one. However, while the sequential scheme always requires J stages, the number of stages for the parallel scheme remains constant for all $J \geq 3$. A key challenge is then to be able to carry out the aggregation stage efficiently. Strategies for aggregation are discussed further in Sect. 4.2.

3.2 Posterior propagation

We will now extend the idea introduced above for arbitrary partitions of \mathbf{Y} and show that this yields an exact decomposition of the full joint distribution (6). As \mathbf{Y} is partitioned along both columns and rows, our hierarchical strategy is conducted in three successive stages. Communication is only required to propagate posteriors from one stage to the next, while within each stage, the subsets are processed in an embarrassingly parallel manner with no communication. The approach, coined *posterior propagation* (PP), proceeds as follows:

Inference stage I Inference is conducted for the parameters of subset $\mathbf{Y}^{(1,1)}$:

$$p(\mathbf{X}^{(1)}, \mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}) \propto p(\mathbf{X}^{(1)}) p(\mathbf{W}^{(1)}) p(\mathbf{Y}^{(1,1)}|\mathbf{X}^{(1)}, \mathbf{W}^{(1)}). \quad (10)$$

Inference stage II Inference is conducted in parallel for parameters of subsets which share columns or rows with $\mathbf{Y}^{(1,1)}$. Posterior marginals from stage 1 are used as priors for the shared parameters:

$$\begin{aligned}
 & p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right) \\
 & \propto p\left(\mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)}\right) p\left(\mathbf{X}^{(i)}\right) p\left(\mathbf{Y}^{(i,1)} | \mathbf{X}^{(i)}, \mathbf{W}^{(1)}\right),
 \end{aligned} \tag{11a}$$

$$\begin{aligned}
 & p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right) \\
 & \propto p\left(\mathbf{X}^{(1)} | \mathbf{Y}^{(1,1)}\right) p\left(\mathbf{W}^{(j)}\right) p\left(\mathbf{Y}^{(1,j)} | \mathbf{X}^{(1)}, \mathbf{W}^{(j)}\right),
 \end{aligned} \tag{11b}$$

for $i = 2, \dots, I$ and $j = 2, \dots, J$.

Inference stage III The remaining subsets are processed in parallel using posterior marginals propagated from stage II as priors:

$$\begin{aligned}
 & p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}, \mathbf{Y}^{(1,j)}, \mathbf{Y}^{(i,j)}\right) \\
 & \propto p\left(\mathbf{X}^{(i)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right) p\left(\mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right) p\left(\mathbf{Y}^{(i,j)} | \mathbf{X}^{(i)}, \mathbf{W}^{(j)}\right).
 \end{aligned} \tag{12}$$

Product density equation Combining the submodels in Eqs. (10–12), for all i and j , and dividing away the multiply-counted propagated posterior marginals yields the following *product density equation*:

$$\begin{aligned}
 & p(\mathbf{X}, \mathbf{W} | \mathbf{Y}) \propto \\
 & p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)}\right) \\
 & \times \prod_{i=2}^I \left[p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right) p\left(\mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)}\right)^{-1} \right] \\
 & \times \prod_{j=2}^J \left[p\left(\mathbf{X}^{(1)}, \mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right) p\left(\mathbf{X}^{(1)} | \mathbf{Y}^{(1,1)}\right)^{-1} \right] \\
 & \times \prod_{i=2}^I \prod_{j=2}^J \left[p\left(\mathbf{X}^{(i)}, \mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}, \mathbf{Y}^{(1,j)}, \mathbf{Y}^{(i,j)}\right) \right. \\
 & \quad \left. \times p\left(\mathbf{X}^{(i)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}\right)^{-1} p\left(\mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}\right)^{-1} \right].
 \end{aligned} \tag{13}$$

The following theorem highlights the fact that this is indeed a proper decomposition of the full posterior density.

Theorem 1 Equation (13) is, up to proportion, an exact decomposition of the full posterior $p(\mathbf{X}, \mathbf{W} | \mathbf{Y})$ given in Eq. (6).

The proof of the theorem is given in Appendix 1.

4 Approximate inference

In the previous section, we introduced a hierarchical decomposition of the joint posterior distribution of the BMF model, which couples inferences over subsets but allows for embarrassingly parallel computations in a sequence of (at most) three stages. The challenge with

implementing this scheme in practice is threefold, and relates to the analytically intractable form of the BMF posterior: (i) propagating posteriors efficiently from one stage to the next, (ii) utilizing the posteriors of one stage as priors in the next stage, and (iii) aggregating all subset posteriors at the end. In this section, we propose to resolve these challenges by using parametric approximations computed from subset posterior samples obtained by MCMC in each stage.

Computational schemes for distributed data settings, combining MCMC with propagation of information through parametric approximations have recently been explored by Xu et al. (2014) and Vehtari et al. (2018). Nevertheless, their expectation propagation algorithms for distributed data require frequent communication among parallel workers to share global information, which could become a bottleneck for large-scale computational problems where the number of model parameters scales linearly with the number of data samples. On the other hand, the proposed method only requires communication between stages of inference. While our focus here is on sampling-based inference, it is worth emphasizing that the decomposition introduced in Sect. 3.2 is itself not tied to any particular inference algorithm. Thus, it could also be combined, e.g. with variational inference.

4.1 Parametric approximations for propagation of posteriors

We present here three alternative approaches for finding tractable approximations from posterior samples. A generic algorithm for the proposed inference scheme using these approximations is given in Algorithm 1.

Algorithm 1: Distributed BMF with approximate posterior propagation.

```

// Inference stage I
1 Draw samples from posterior of BMF with data  $\mathbf{Y}^{(1,1)}$  and priors  $p(\mathbf{X}), p(\mathbf{W})$ ;
2 Find approximations  $\hat{p}(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}), \hat{p}(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)})$ ;
// Inference stage II
3 for  $i = 2 : I$  in parallel do
4   Draw samples from posterior of BMF with data  $\mathbf{Y}^{(i,1)}$  and priors  $p(\mathbf{X}), \hat{p}(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)})$ ;
5   Find approximations  $\hat{p}(\mathbf{X}^{(i)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}), \hat{p}(\mathbf{W}^{(1)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)})$ ;
6 end
7 for  $j = 2 : J$  in parallel do
8   Draw samples from posterior of BMF with data  $\mathbf{Y}^{(1,j)}$  and priors  $\hat{p}(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}), p(\mathbf{W})$ ;
9   Find approximations  $\hat{p}(\mathbf{X}^{(1)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}), \hat{p}(\mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)})$ ;
10 end
// Inference stage III
11 for  $i = 2 : I$  in parallel do
12   for  $j = 2 : J$  in parallel do
13     Draw samples from posterior of BMF with data  $\mathbf{Y}^{(i,j)}$  and priors  $\hat{p}(\mathbf{X}^{(i)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}),$ 
       $\hat{p}(\mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)})$ ;
14     Find approximations  $\hat{p}(\mathbf{X}^{(i)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}, \mathbf{Y}^{(i,j)}), \hat{p}(\mathbf{W}^{(j)}|\mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}, \mathbf{Y}^{(i,j)})$ ;
15   end
16 end
// Aggregation
17 Aggregate subset posteriors to approximate the full-data posterior using Algorithm 2;
```

Gaussian mixture model approximation For the first approach, we note that the posterior distributions represented by the samples are typically multimodal due to the inherent unidentifiability of the BMF model. *Gaussian mixture models* (GMM) are universal approximators of probability distributions, that is, given sufficiently many components, they can approximate any continuous distribution with arbitrary accuracy. Thus, they are a reasonable parametric approximation of the posterior.

Dominant mode approximation Our second approach is based on the intuition that for purposes of prediction in matrix completion tasks, it is sufficient to find only one of the possibly infinitely many solutions to the matrix factorization problem. In this approach, we therefore locate the *dominant mode* from each posterior distribution. We then fit a multivariate Gaussian to the samples corresponding to this mode only, and propagate it as a prior to the following stage.

Moment matching approximation Our final approach presents an intermediate between the previous two approaches. Here, we fit a unimodal multivariate Gaussian to the entire set of posterior samples for each parameter using *moment matching*. Beyond its simplicity, propagating Gaussian approximations for priors has the appeal that the inferences in different stages (i.e. steps 1, 4, 8, 13 in Algorithm 1) can be processed as a standard BMF. It also has the usual interpretation that the log-posterior corresponds to a sum-of-squared-errors objective function with quadratic regularization terms (Salakhutdinov and Mnih 2008b). Finally, the moment matching approximation brings our scheme in close relation to recent work on expectation propagation for distributed data (Xu et al. 2014; Vehtari et al. 2018), but with only limited-communication and a single pass over the data as in assumed density filtering.

4.2 Approximating the product density equation for aggregation

Each subset posterior inference results in a joint distribution for subsets of the parameters \mathbf{X} and \mathbf{W} , approximated by a set of samples. Direct aggregation of these joint distributions using the product density equation (13) is a computationally challenging task. For computational efficiency, and to enable the use of the approximations introduced above, we simplify the task by decoupling the parameters and performing the aggregation by posterior marginals. With parametric approximations for each subset posterior marginal available (steps 2, 5, 9, 14 in Algorithm 1), we aggregate them by multiplying them together and dividing away all multiply counted propagated posteriors.

We assume that the marginal distributions over the parameter matrices can be factorized along rows into a product of K -dimensional distributions, i.e.

$$\hat{p}(\mathbf{X}|\mathbf{Y}) = \prod_{n=1}^N \hat{p}(\mathbf{x}_n|\mathbf{Y}), \quad \hat{p}(\mathbf{W}|\mathbf{Y}) = \prod_{d=1}^D \hat{p}(\mathbf{w}_d|\mathbf{Y}).$$

The *dominant mode* and *moment matching* approximations both produce unimodal multivariate Gaussian representations for each row of the parameter matrices. By the properties of Gaussian distributions, the aggregated posterior for the n th row of \mathbf{X} is then obtained as

$$\begin{aligned}
\hat{p}(\mathbf{x}_n|\mathbf{Y}) &= \mathcal{N}_K\left(\mathbf{x}_n \mid \hat{\mu}_{\mathbf{x}_n}^*, [\hat{\Lambda}_{\mathbf{x}_n}^*]^{-1}\right), \\
\hat{\Lambda}_{\mathbf{x}_n}^* &= \hat{\Lambda}_{\mathbf{x}_n}^{(1)} + \sum_{j=2}^J \left(\hat{\Lambda}_{\mathbf{x}_n}^{(j)} - \hat{\Lambda}_{\mathbf{x}_n}^{(1)}\right), \\
\hat{\mu}_{\mathbf{x}_n}^* &= [\hat{\Lambda}_{\mathbf{x}_n}^*]^{-1} \left(\hat{\Lambda}_{\mathbf{x}_n}^{(1)} \hat{\mu}_{\mathbf{x}_n}^{(1)} + \sum_{j=2}^J \left(\hat{\Lambda}_{\mathbf{x}_n}^{(j)} \hat{\mu}_{\mathbf{x}_n}^{(j)} - \hat{\Lambda}_{\mathbf{x}_n}^{(1)} \hat{\mu}_{\mathbf{x}_n}^{(1)}\right) \right), \quad (14)
\end{aligned}$$

where $\hat{\mu}_{\mathbf{x}_n}^{(j)}, \hat{\Lambda}_{\mathbf{x}_n}^{(j)}, j = 1, \dots, J$, denote the estimated statistics of the posterior for the j th submodel. Note that for submodels indexed by $j = 2, \dots, J$, the effect of the first submodel ($j = 1$) has been removed. The aggregation of each $\hat{p}(\mathbf{w}_d|\mathbf{Y})$ is done in similar fashion.

For the GMM approach, the posterior marginal for each \mathbf{w}_d and \mathbf{x}_n is a mixture with density $f(x) = \sum_c \hat{\pi}_c \cdot \mathcal{N}(\mathbf{x}_n; \hat{\mu}_{\mathbf{x}_n}^c, [\hat{\Lambda}_{\mathbf{x}_n}^c]^{-1})$. Maintaining this approximation in the aggregation phase would lead to the computationally challenging problem of dividing one mixture by another. Emphasizing speed and efficiency, we instead apply Eq. (14) using pooled mixture components:

$$\begin{aligned}
\hat{\mu}_{\mathbf{x}_n}^{(j)} &= \sum_{c=1}^C \hat{\pi}_c \cdot \hat{\mu}_{\mathbf{x}_n}^c \\
[\hat{\Lambda}_{\mathbf{x}_n}^{(j)}]^{-1} &= \sum_{c=1}^C \left(\hat{\pi}_c [\hat{\Lambda}_{\mathbf{x}_n}^c]^{-1} + \hat{\pi}_c \cdot (\hat{\mu}_{\mathbf{x}_n}^c - \hat{\mu}_{\mathbf{x}_n}^{(j)}) (\hat{\mu}_{\mathbf{x}_n}^c - \hat{\mu}_{\mathbf{x}_n}^{(j)})^\top \right).
\end{aligned}$$

To improve the numerical stability of using Eq. (14), we additionally apply an eigenvalue correction to correct for occasionally occurring non-positive definite matrices in the aggregation, which is summarized in Algorithm 2.

Algorithm 2: Posterior marginal aggregation of $p(\mathbf{x}_n|\mathbf{Y})$ using eigenvalue correction. The function `EigenvalueCorrection(M)` finds the smallest eigenvalue of a non-positive definite symmetric matrix \mathbf{M} , and adds its absolute value and a small constant to all diagonal elements. The aggregation of $p(\mathbf{w}_d|\mathbf{Y})$ is done analogously.

```

1 for  $j = 2 : J$  in parallel do
2   if  $\hat{\Lambda}_{\mathbf{x}_n}^{(j\setminus 1)} = \hat{\Lambda}_{\mathbf{x}_n}^{(j)} - \hat{\Lambda}_{\mathbf{x}_n}^{(1)}$  is positive definite then
3      $\hat{\Lambda}_{\mathbf{x}_n}^{*(j)} = \hat{\Lambda}_{\mathbf{x}_n}^{(j)}$ 
4   else
5      $\hat{\Lambda}_{\mathbf{x}_n}^{*(j\setminus 1)} = \text{EigenvalueCorrection}(\hat{\Lambda}_{\mathbf{x}_n}^{(j\setminus 1)});$ 
6      $\hat{\Lambda}_{\mathbf{x}_n}^{*(j)} = \hat{\Lambda}_{\mathbf{x}_n}^{*(j\setminus 1)} + \hat{\Lambda}_{\mathbf{x}_n}^{(1)}$ 
7   end
8 end
9  $\hat{\Lambda}_{\mathbf{x}_n}^* = (2 - J) \hat{\Lambda}_{\mathbf{x}_n}^{(1)} + \sum_{j=2}^J \hat{\Lambda}_{\mathbf{x}_n}^{*(j)};$ 
10  $\hat{\mu}_{\mathbf{x}_n}^* = [\hat{\Lambda}_{\mathbf{x}_n}^*]^{-1} \left( (2 - J) \hat{\Lambda}_{\mathbf{x}_n}^{(1)} \hat{\mu}_{\mathbf{x}_n}^{(1)} + \sum_{j=2}^J \hat{\Lambda}_{\mathbf{x}_n}^{*(j)} \hat{\mu}_{\mathbf{x}_n}^{(j)} \right);$ 

```

4.3 Scalability

This section provides a brief discussion about the scalability of the above inference scheme in terms of computation time and communication cost. With U workers available, both rows and columns can be partitioned into $\sqrt{U} + 1$ parts, assuming for simplicity an equal number of partitions in both directions (note, however, that this is not a requirement for our method). This results in a total of $U + 2\sqrt{U} + 1$ subsets. The computational cost of a typical BMF inference algorithm per iteration is proportional to $(N + D)K^3 + MK^2$, where N and D are the respective dimensions of the observation matrix, M is the number of observed values, and K is the number of latent dimensions. Thus, for each submodel, the theoretical computation time is proportional to

$$t_0 := \left[(N + D)K^3/(\sqrt{U} + 1) + MK^2/(U + 2\sqrt{U} + 1) \right] T,$$

assuming an equal number of observations in each subset and T iterations. Thus, the initial stage can be completed with one worker in time t_0 , inference stage II can be processed with $2\sqrt{U}$ workers in time t_0 , and inference stage III can be completed with U workers in time t_0 . Finally, the aggregation step mainly involves calculating the product of multivariate Gaussian distributions, which can be done with $2(\sqrt{U} + 1)$ parallel workers in time proportional to

$$t_a := \frac{\max(N, D)}{\sqrt{U} + 1} (K + K^2).$$

Therefore, the total computation time of the algorithm with U worker nodes is proportional to the sum of the computation times of each inference stage plus the computation time of the aggregation, $t = 3t_0 + t_a$.

In terms of communication cost, the proposed inference scheme requires first communicating inputs to workers and then collecting the outputs for aggregation. The inputs consist of two parts: data and prior distributions. As workers use non-overlapping parts of the data, the total amount of communication needed to load the data is proportional to the number of observations M . Each worker receives parameters for $(N + D)/(\sqrt{U} + 1)$ distributions, each with L parameters; for the dominant mode and moment matching approximations L is proportional to $K + K^2$ and for the Gaussian mixture model approximation it is proportional to $C(K + K^2)$, where C is the number of components and K^2 is due to using a full covariance matrix for the posterior of parameters. As there are U workers, the total amount of communication needed for input distributions is proportional to $\sqrt{U}(N + D)L$. The output distributions are of the same size as the input distributions. Thus, the communication cost at the aggregation stage is the same as the communication cost of input distributions.

5 Experiments

In this section, we evaluate the empirical performance of our limited-communication inference scheme for BMF, **posterior propagation with parametric approximations**, by comparing it with both embarrassingly parallel MCMC methods in Sect. 5.4, and available state-of-the-art parallel and distributed implementations of BMF and non-negative matrix factorization (NMF) in Sect. 5.5. In Sect. 5.6, we further analyse the advantage of our method over embarrassingly parallel MCMC in terms of encouraging a common representation for model parameters to facilitate subset posterior aggregation. Details about the implementation, test setup and data sets are provided in Sects. 5.1–5.3.

5.1 Implementation

For posterior inference in each subset, we use our R implementation² of the BPMF Gibbs sampler presented by Salakhutdinov and Mnih (2008a).³ BPMF considers the noise precision τ as a constant and places a normal-Wishart prior on the hyperparameters $\mu_{\mathbf{W}}$ and $\Lambda_{\mathbf{W}}$, as well as on hyperparameters $\mu_{\mathbf{X}}$ and $\Lambda_{\mathbf{X}}$. In the first stage of PP, we sample all parameters of the BPMF model. However, in the second and third stages, we sample hyperparameters $\mu_{\mathbf{W}}$ and $\Lambda_{\mathbf{W}}$ only when the posterior of \mathbf{W} is not propagated. Similarly, hyperparameters $\mu_{\mathbf{X}}$ and $\Lambda_{\mathbf{X}}$ are sampled only when the posterior of \mathbf{X} is not propagated.

We have introduced three alternative approaches (in Sect. 4.1) to estimate subset posteriors from samples:

1. The GMM approximation fits a mixture of multivariate Gaussians to a clustering of the posterior samples (**PP GMM**).
2. The dominant mode approximation fits a multivariate Gaussian to the samples of the dominant mode (**PP DM**).
3. The moment matching approximation fits a unimodal multivariate Gaussian to the entire set of posterior samples (**PP MM**).

For a computationally fast way of implementing the first two algorithms, we first use the nonparametric λ -means clustering algorithm (Comiter et al. 2016) to cluster the posterior samples, then (i) for **PP DM** we choose the cluster with the maximum number of posterior samples to estimate the posterior, (ii) for **PP GMM** we estimate the posteriors for top- N modes/clusters. In our experiments, we use the top-3 modes. When using the estimated GMM as a prior for BMF, we perform Gibbs sampling in a similar way as for mixture models; denoting by $\hat{\mu}_{\mathbf{x}_n}^c$ and $\hat{\Lambda}_{\mathbf{x}_n}^c$ the estimated mean and precision, respectively, of mode c in the posterior of parameter \mathbf{x}_n :

1. Compute the probability $p(\mathbf{x}_n | \hat{\mu}_{\mathbf{x}_n}^c, [\hat{\Lambda}_{\mathbf{x}_n}^c]^{-1})$ of generating the parameter for each mixture component, i.e. the likelihood of \mathbf{x}_n .
2. Calculate the responsibility $\gamma(\mathbf{x}_n^c) = \hat{\pi}_c \cdot p(\mathbf{x}_n | \hat{\mu}_{\mathbf{x}_n}^c, [\hat{\Lambda}_{\mathbf{x}_n}^c]^{-1})$ of each component to explain \mathbf{x}_n .
3. Choose the component with the maximum $\gamma(\mathbf{x}_n^c)$ as the propagated prior for \mathbf{x}_n , and update the parameter with its statistics.

The above procedure is done analogously for \mathbf{w}_d .

5.2 Test setup

We evaluate the distributed inference methods using simulated data and three real-world data sets: **MovieLens-1M**, **MovieLens-20M** and **ChEMBL**. In addition to inference on full data for medium-sized data sets, we predict missing values using column means; this benchmark serves as a baseline and sanity check.

² An implementation of BMF with PP, with highly optimized C libraries, is available in SMURFF software on github (bmffpp branch): <https://github.com/ExaScience/smurff>.

³ We have corrected an apparent error in the original formulation of the update formula for $[\mathbf{W}_0^*]^{-1}$, where \bar{S} should be calculated with $\bar{S} = \frac{1}{N} \sum_{i=1}^N (\mathbf{x}_i - \bar{\mathbf{X}})(\mathbf{x}_i - \bar{\mathbf{X}})^T$ (Murphy 2007) rather than $\bar{S} = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i \mathbf{x}_i^T$. Alternatively, one could use the original formula for \bar{S} and calculate $[\mathbf{W}_0^*]^{-1}$ with $[\mathbf{W}_0^*]^{-1} = \mathbf{W}_0^{-1} + N\bar{S} + \frac{\beta_0 N}{\beta_0 + N}(\mu_0 - \bar{\mathbf{X}})(\mu_0 - \bar{\mathbf{X}})^T - N\bar{\mathbf{X}}(\bar{\mathbf{X}})^T$ or $[\mathbf{W}_0^*]^{-1} = \mathbf{W}_0^{-1} + N\bar{S} + \beta_0 \mu_0 (\mu_0)^T - \beta_0^* \mu_0^* (\mu_0^*)^T$ (Teh 2007).

We evaluate performance by predictive accuracy. To this end, we randomly partition the data into training and test sets and use root mean squared error (RMSE) on the test set as performance measure. For prediction in the experiments, we use Eq. (1) with posterior means as values for \mathbf{X} and \mathbf{W} . Furthermore, we use wall-clock time⁴ to measure the speed-up achieved by parallelization. The reported wall-clock time for our method is calculated by summing the maximum wall-clock times of submodels for each inference stage plus the wall-clock time of the aggregation step.⁵

In all experiments, we ran Gibbs sampling with 1200 iterations. We discarded the first 800 samples as burn-in and saved every second of the remaining samples yielding in total 200 posterior samples. The results were averaged over 5 runs. For parallelization, we experiment with different partitioning schemes; a partitioning scheme $r \times c$ means that rows are partitioned into r and columns into c subsets. The partitioning scheme 1×1 refers to the full data. Note that the maximum number of parallel workers that can be used by our algorithm is equal to the number of subsets in the third stage, i.e. $(r - 1) \times (c - 1)$. We also tested two ordering schemes. In the first scheme, rows and columns are permuted randomly (**row/column order: random**). In the second scheme, the rows and columns are reordered into a descending order according to the proportion of observations in them (**row/column order: decreasing**). Thus, the most dense rows and columns are processed in the first two stages, by which the subsequently propagated posteriors can be made more informative.

The configuration of compute nodes that we used to run the experiments in Sects. 5.4 and 5.5 is given in Appendix 2.

5.3 Data sets

The **MovieLens-1M** (Harper and Konstan 2015) data set consists of 1,000,209 movie ratings by 6040 users on 3706 movies. Approximately 4.5% of the elements of the movie rating matrix, where each user corresponds to a row and each movie to a column, are observed. The **MovieLens-20M** (Harper and Konstan 2015) data set contains 20 million ratings from 138,493 users on 27,278 movies; that is, about 0.53% of the elements are observed. Following Simm et al. (2015), we set $\tau = 1.5$ and $K = 10$ for performance analysis.

The **ChEMBL** (Bento et al. 2014) data set describes interactions between drugs and proteins using the IC50 measure. The data set has 15,703 rows corresponding to drugs and 346 columns corresponding to proteins, and contains 59,280 observations which is slightly over 1% of the elements. Again, we follow Simm et al. (2015) to set $\tau = 5$ and $K = 10$. As ChEMBL contains only 346 columns, we only partitioned the rows.

For these real-world data sets, we conduct a fivefold cross-validation study where 20% of observations are withheld and used as test set.

To complement the real data sets, we generated **simulated data** sets with 6040 observations and 3706 features as follows: We set the number of latent factors to $K = 5$. The elements of the matrices \mathbf{W} and \mathbf{X} were generated independently from the standard univariate normal distribution. Finally, we generated the data with $\mathbf{Y} = \mathbf{XW}^\top + \boldsymbol{\epsilon}$, where the $\boldsymbol{\epsilon}$ is a noise matrix whose elements were generated from a standard normal distribution. For

⁴ Wall-clock time measures the real time between the start and the end of a program. For parallel processes, we use the wall-clock time of the slowest process.

⁵ In our current implementation, we run each stage on a cluster as an array job consisting of multiple independent processes. While in this implementation communication between stages is done using read and write operations on disk, its effect on total computation time is negligible as this needs to be done at most three times during the entire program, and the amount of data communicated is relatively small.

learning, we set the parameters K and τ to the corresponding correct values, i.e., $K = 5$ and $\tau = 1$. We generated 5 independent simulated data sets.

In many real-world applications, such as collaborative filtering and the ChEMBL benchmark, the data are very sparsely observed. We analyse the predictive performance of the model with respect to different types of missing data. To this end, we randomly select 80% of the data as missing, use these missing data as test set and the remaining data as training set. To additionally simulate not-missing-at-random data as the second simulated data scenario, we first assigned weights w_n and w_d to each row and column, respectively, such that they form an equally spaced decreasing sequence $0.9, \dots, 0.005$. Then we assigned the element y_{nd} to the test data with probability $w_n w_d$; this results in a matrix with about 80% of elements missing. This is referred to as the structured missingness scenario.

5.4 Comparison with embarrassingly parallel methods

In this subsection, we compare the predictive performance and computation times of the proposed inference scheme to those of the full model, as well as the following algorithms for embarrassingly parallel MCMC⁶:

1. **Parametric density product** (parametric) is a multiplication of Laplacian approximations to subset posteriors. The aggregated samples are drawn from the resulting multivariate Gaussian (Neiswanger et al. 2014).
2. **Semiparametric density product** (semiparametric) draws aggregated samples from multiplied semi-parametric estimates of subset posteriors (Neiswanger et al. 2014).
3. **Random partition tree** (randomPARTree) works by first performing space partitioning over subset posteriors, followed by a density aggregation step which simply multiplies densities across subsets for each block and then normalizes (Wang et al. 2015). Aggregated samples are drawn from the aggregated posterior.

All of the above algorithms are implemented in the Matlab PART library.⁷ We ran the randomPARTree algorithm with different values for its hyperparameters (i.e. min cut length=0.001 or 0.01, min fraction block = 0.01 or 0.1, cut type = “kd” or “ml”, local gaussian smoothing = true or false) for pilot analysis, and found that there are no significant differences in the predictive performance for different values for the hyperparameters. Thus, for this algorithm, we use KD-tree for space partition and 1000 resamples for final approximation, and use the default values for the other hyperparameters provided in the library.

5.4.1 Results

The results for ChEMBL, MovieLens-1M, MovieLens-20M, and simulated data are shown in Figs. 2 and 3. To improve the readability of the plots, we only plot RMSE for the two posterior aggregation methods that give the best performance for embarrassingly parallel MCMC on each data set. In the following, we summarize the major conclusions of this evaluation.

As a general conclusion, we found that posterior propagation can give almost an order of magnitude faster computation times with a negligible effect on predictive accuracy, compared to MCMC inference on the full data matrix; this can be seen on simulated data and

⁶ We found that running BMF with Gibbs sampler for each subset using down-weighted priors $p(\mathbf{X})^{1/I}$ and $p(\mathbf{W})^{1/J}$ could lead to numerical instabilities (i.e. resulting in non-positive definite precision matrices). Therefore, we chose to run BMF for each subset with the standard normal priors $p(\mathbf{X})$ and $p(\mathbf{W})$, and then divided away the multiply-counted priors when aggregating the subset posteriors.

⁷ <https://github.com/wwrechard/random-tree-parallel-MCMC>.

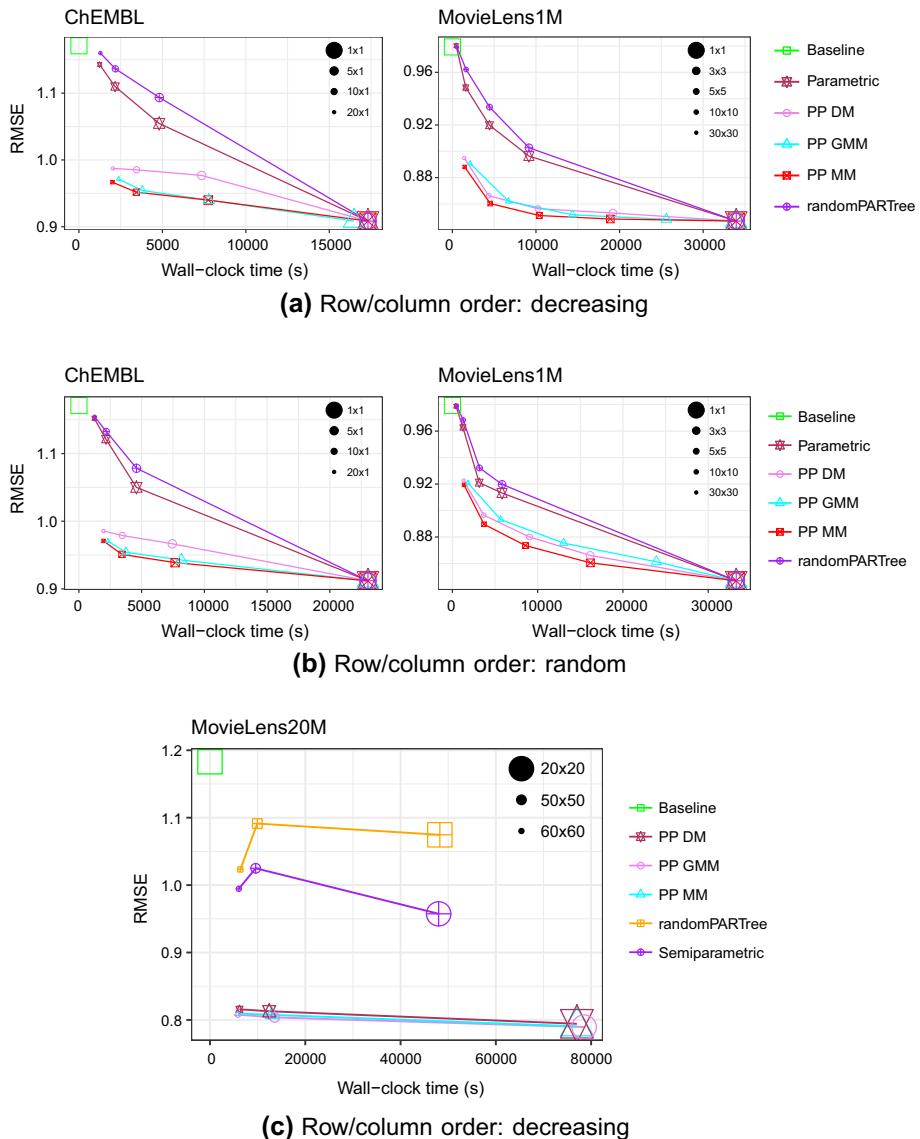


Fig. 2 Test RMSE and wall-clock time on ChEMBL (left) and MovieLens-1M (right) data in **a**, **b**, and MovieLens-20M data in **c** with $K = 10$. The size of legends/symbols indicates different partition schemes. Lower RMSE is better. **PP MM** works best for all three data sets, followed by **PP GMM** and **PP DM**

MovieLens (Fig. 3 and right-hand side of Fig. 2a). The almost horizontal RMSE curves for our methods on MovieLens-20M and simulated data indicate that posterior propagation speeds up computation with almost no loss on accuracy. Note that without approximations, PP would give the same results as the full model. The difference between them therefore *quantifies the effect of the approximations* made in our approach.

Of the embarrassingly parallel MCMC methods, the parametric aggregation method gives the best predictive accuracy on ChEMBL and MovieLens-1M data. Posterior propagation

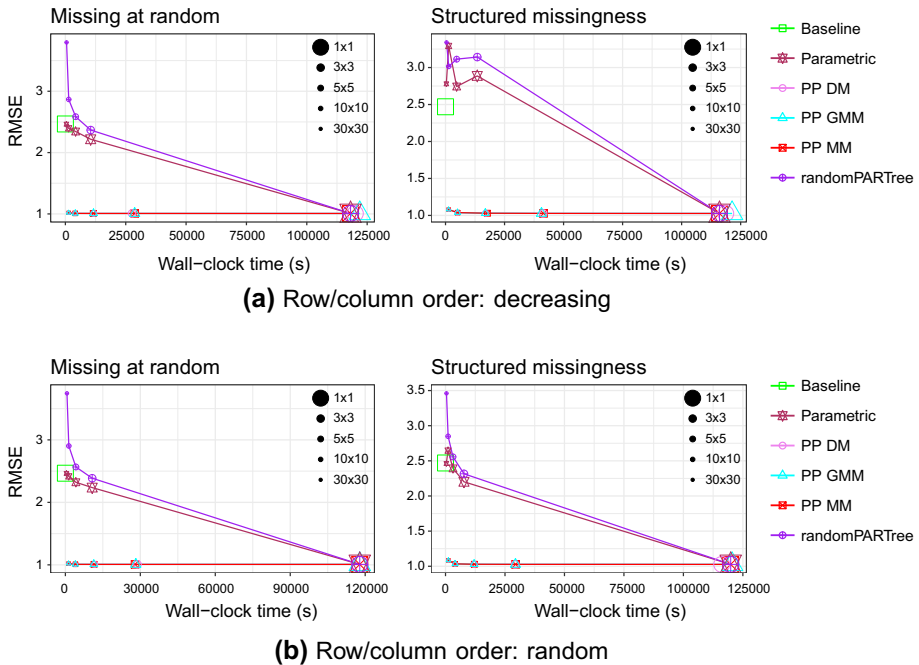


Fig. 3 Test RMSE and wall-clock time on simulated data with 80% missing values. Lower RMSE is better. Left: values missing at random. Right: missing data generated with the structured missingness scenario. Nearly horizontal curves imply that posterior propagation speeds up computation with almost no loss on accuracy

provides better predictive accuracy (lower RMSE values) than any of the embarrassingly parallel MCMC methods on all of the data sets considered. We also found out that reordering rows and columns, in a decreasing order with respect to the number of observations, usually improves the accuracy of posterior propagation compared to using a random order of rows and columns; this can be seen on the sparsely observed MovieLens-1M data (right-hand side of Fig. 2 (a,b)). In Appendix 3, we analyse the results in Figs. 2 and 3 from another perspective to show the wall-clock time speed-up as a function of the number of parallel workers.

We further explored empirically whether posterior propagation can produce good prediction for users and items with only a few observations. This is useful for cold-start problems, i.e., recommendation for new users with very few observed ratings. For this analysis, we visualize test RMSE versus the number of ratings per user in the training set in Fig. 4. Again, we observed that compared with the alternative embarrassingly parallel MCMC methods, our methods show superior performance for all user groups and improve prediction for users with very few observed ratings.

5.5 Comparison with other parallel and distributed implementations

In this subsection, we show that our method achieves competitive results compared to alternative implementations of parallel and distributed BMF, while keeping the communication requirement bounded. To this end, we compare our method on large-scale data (MovieLens-20M) with the following implementations:

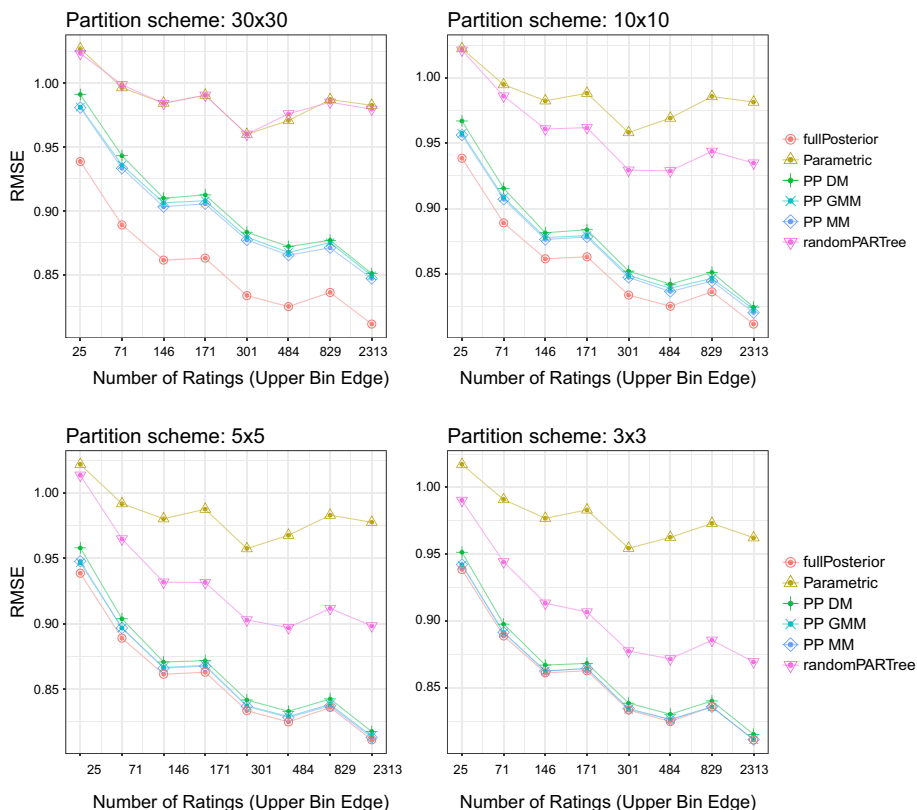


Fig. 4 Test RMSE with respect to user frequency (i.e. the number of ratings the users have in the training set) on MovieLens-1M (row/column order: decreasing) data with $K = 10$ for different partition schemes. Here the fullPosterior refers to the posterior on the full data (i.e. partition scheme 1×1). Compared with the embarrassingly parallel methods, our proposed methods can produce predictions with lower variance for all user groups and produce predictions comparable with that of the full posterior if the size of the subset is large enough (e.g. subset with 1208 rows by 742 columns for partition scheme 5×5)

1. **Distributed parallel BPMF**⁸ (D-BPMF): a state-of-the-art C++ implementation of distributed BPMF with Gibbs sampler (Vander Aa et al. 2017). It supports hybrid communication for distributed learning, which utilizes TBB for shared memory level parallelism and Global Address Space Programming Interface (GASPI) for cross-node communication.
2. **NMF with parallel SGLD**⁹ (NMF + PSGLD): OpenMP implementation of non-negative matrix factorization with parallel SGLD (Şimşekli et al. 2017). This is an open source software that is similar to BPMF with distributed SGLD¹⁰ (Ahn et al. 2015). This software requires careful tuning of hyperparameters in order to avoid numerical instabilities/overflow issues and get reasonable predictions. We set $\epsilon = 0.0001$, $\beta = 2$ (using a Gaussian likelihood), $\lambda = 0.01$, $initStd = 0.5$ for the experiment based on a pilot study.

⁸ <https://github.com/ExaScience/bpmf>.

⁹ https://perso.telecom-paristech.fr/simsekli/nmf_sgmmc/.

¹⁰ The software for large scale BPMF with distributed SGLD is not publicly available.

Table 2 Comparison of performance for different implementations of distributed BMF/NMF on MovieLens-20M data

Model	#Nodes/#cores	RMSE	Wall-clock time (s)	Communication cost (floating-point unit)
Proposed method BMF+PP	361/8664	0.789 ± 0.0003	70	$3.15 \times 10^6 \cdot (K + K^2)$
D-BPMF (Vander Aa et al. 2017)	1/24	0.779 ± 0.0003	204	–
	2/48	0.781 ± 0.001	158	$1.65 \times 10^5 \cdot K \cdot T$
	8/192	0.780 ± 0.0051	89	$1.16 \times 10^6 \cdot K \cdot T$
	64/1536	0.772 ± 0.0074	124	$1.04 \times 10^7 \cdot K \cdot T$
	85/2040	0.780 ± 0.0182	120	$1.39 \times 10^7 \cdot K \cdot T$
NMF + PSGLD (Şimşekli et al. 2017)	1/24	0.903 ± 0.0006	9837	–

The communication costs are computed using the formulas in Table 1 and consider cross-node (distributed memory) communication only. Note that for D-BPMF, communication takes place in every iteration, whereas for the proposed method communication is only required between stages of inference. Also note that in practice the factor $(K + K^2)$ in the cost of the proposed method reduces to $K(K + 3)/2$ due to the precision matrix being symmetric

Bold values indicate the minimum averaged RMSE over the 5 runs

For our method, we used a 20×20 partition scheme with the same setup as for the experiments in Fig. 2c: i.e. $K = 10$, $T = 1200$ iterations for Gibbs sampling and a burn-in of 800 samples. Note that our method was implemented in the R language without any optimization, while the other two methods were implemented with highly optimized C libraries. For the sake of obtaining comparable results, the RMSE was obtained using our R implementation (same as in Fig. 2c), while the wall-clock time is an estimate computed by using D-BPMF within each individual data subset in the three stages of our posterior propagation scheme. For each subset, we used a single node with 24 cores, resulting in a total of $(20 - 1)^2 = 361$ parallel processes (one per node) for the third stage.

Since only the OpenMP implementation of NMF + PSGLD was available to us, it was run within a single compute node with no cross-node communication.

5.5.1 Results

The results of the comparison are shown in Table 2. The approximations made in our approach (BMF + PP) have only a slight negative effect on the RMSE compared to D-BPMF, which makes no distributional approximations and should therefore be close to the accuracy of the full model (not computed here because of its large size). On the other hand, in terms of computation time, our method is able to leverage the combination of a high level of parallelism and a very low communication frequency compared to D-BPMF, which requires frequent communication. Varying the number of nodes/cores for the latter yields results which are consistent with the empirical finding of Vander Aa et al. (2017): parallel efficiency initially improves with increased parallelism, but begins to level off as the number of nodes increases beyond the boundary of fast network connections in the HPC cluster. A further point to note is that the communication cost of D-BPMF increases linearly with respect to the number of MCMC samples while ours stays constant, thus, the longer chains we run, the more advantage

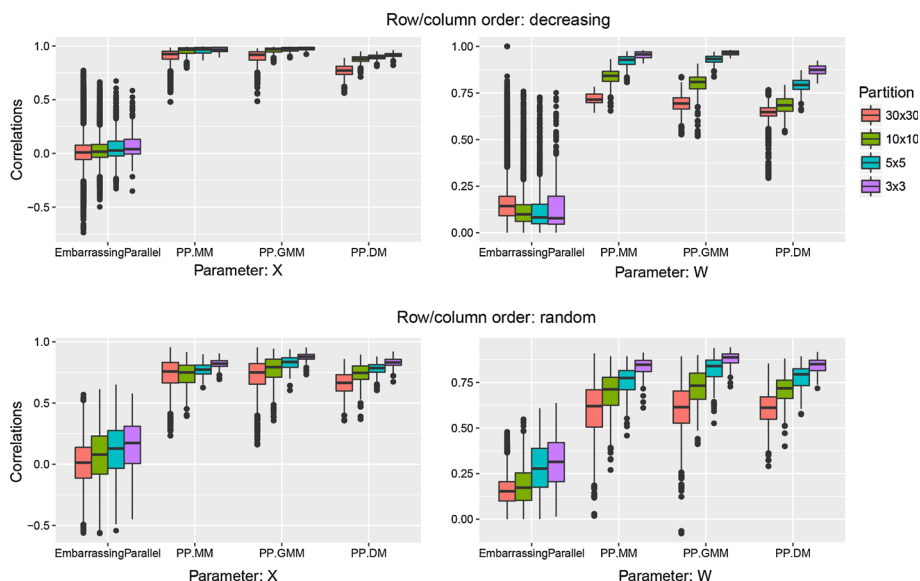


Fig. 5 The correlations of posterior estimates of different subsets for different variants of the proposed method and embarrassingly parallel MCMC on MovieLens-1M data with $K = 10$, for different partition schemes. Compared to embarrassingly parallel MCMC, the proposed method can produce posterior estimates which are highly correlated between different subsets, suggesting that it can enforce a common representation for model parameters, making the aggregation of submodels feasible

we get. Finally, NMF + PSGLD performs worse than the other alternative methods in terms of predictive accuracy and wall-clock computation time. This is partially due to the difficulty of tuning the hyperparameters for each specific data set for the DSGLD and PSGLD methods.

5.6 Correlations of the subset posteriors

We observed in Sect. 5.4 that compared with existing embarrassingly parallel methods, the proposed method can provide a better trade-off between predictive accuracy and wall-clock time on several benchmark data sets. In this section, we investigate empirically to what extent our method is able to encourage joint identifiability via dependencies between the inferences in different subsets.

For this purpose, we compute for our method the correlations between the posterior expected means of parameters in subsets sharing rows or columns, and compare them to the corresponding correlations produced by embarrassingly parallel MCMC, see Fig. 5. For example, for the partition scheme in Fig. 1, we would calculate the correlations of posterior means for subsets as follows: $\text{cor}(\hat{\mathbf{X}}^{(1,1)}, \hat{\mathbf{X}}^{(1,j)})$, $\text{cor}(\hat{\mathbf{X}}^{(i,1)}, \hat{\mathbf{X}}^{(i,j)})$, $\text{cor}(\hat{\mathbf{W}}^{(1,1)}, \hat{\mathbf{W}}^{(i,1)})$, $\text{cor}(\hat{\mathbf{W}}^{(1,j)}, \hat{\mathbf{W}}^{(i,j)})$, for $i = 2, \dots, I$, $j = 2, \dots, J$. For embarrassingly parallel MCMC, to avoid low correlations due to mis-aligned permutations of the latent dimensions in different submodels, highly correlated latent dimensions were aligned prior to calculating the correlations between the posterior means.

An obvious trend from Fig. 5 is that the correlation scores of posterior estimates generated by our method are much higher than those of embarrassingly parallel MCMC. The observation suggests that by propagating the posteriors obtained from the earlier stage to the next stage as

priors, our method can produce highly dependent subset posteriors. On the other hand, since existing embarrassingly parallel MCMC methods do not introduce any dependencies between the inferences for different subsets, they are unable to enforce a common permutation and scaling for parameters, making the aggregation step challenging for unidentifiable models.

6 Discussion

We have introduced a hierarchical embarrassingly parallel strategy for Bayesian matrix factorization, which enables a trade-off between accuracy and computation time, and uses very limited communication. The empirical evaluation on both real and simulated data shows that (i) our distributed approach is able to achieve a speed-up of almost an order-of-magnitude, with a negligible effect on predictive accuracy, compared to MCMC inference on the full data matrix; (ii) our method also significantly outperforms state-of-the-art embarrassingly parallel MCMC methods in accuracy, and (iii) performs comparably to other available distributed and parallel implementations of BMF. We further show that, unlike existing embarrassingly parallel approaches, our method produces posterior estimates, which are highly correlated across different subsets and thus enable a meaningful aggregation of subset inferences.

We have experimented with both inclusive approximations (GMM and MM; attempting to include all sampled modes, which is still restricted to a small finite number) as well as exclusive approximations (DM; attempting to exclude all but one mode, a property shared with variational inference). In our current setting, the inclusive approximations gave more consistent performance, striking a balance between restricting the set of solutions to encourage identifiability and letting the sampler explore good solutions. While the proposed approximations work well, more accurate representations for subset posteriors could be considered instead, in particular for aggregation (e.g. Wang et al. 2015). This would be relevant especially if we were interested in an accurate, global representation of the joint distribution of \mathbf{X} and \mathbf{W} , instead of predicting unobserved elements of the data matrix \mathbf{Y} , which was the case in our current work. While more sophisticated representations may improve accuracy, they come at the expense of increased computational burden. An additional motivation to consider alternative subset posterior representations is to be able to establish theoretical guarantees for the global model, which despite their good empirical performance, are not straightforward to give for the current approximations.

Several works on distributed learning for BMF [e.g. BMF with DSGLD (Ahn et al. 2015), NMF with parallel SGLD (Şimşekli et al. 2017)] assume (implicitly) that the models are trained with blocks in an orthogonal group (or squared partition), in order to avoid conflicting access to parameters among parallel workers. In our work, we do not make any assumptions about the partition scheme, and our method can therefore work flexibly with diversified partition schemes, which depend on the size of the data in both dimensions. For instance, it can work with partitions only along the row direction for tall data, or partitions only along the column direction for fat data, or partitions along both row and column directions for tall and fat data.

The focus of our work has been to develop an efficient and scalable distributed learning scheme for BMF. In doing so, we have assumed implicitly that the data are missing at random, following many other works on BMF. While many (if not most) real-world data sets exhibit non-random patterns on missingness, we have handled such patterns using the simple strategy of reordering rows and columns into a descending order according to the proportion of observations in them. Thus, the most dense rows and columns are processed

during the first two stages, by which the subsequently propagated posteriors can be made more informative. However, it is also possible to handle non-random patterns of missing values in a more principled manner. In the context of matrix factorization, Hernández-Lobato et al. (2014) modelled the generative process for both the data and the missing data mechanism, and showed empirically that learning these two models jointly can improve performance of the MF model. This strategy would be straightforward to incorporate within our distributed scheme.

Finally, we have run experiments on a scale of only tens of millions of elements, but there is no obstacle for running the proposed distributed algorithm on larger matrices. Indeed, the proposed approach as such is not implementation-dependent, and it could be used together with any available well-optimized (but more communication intensive) implementation to enable further scaling beyond the point at which parallel efficiency would otherwise begin to level off.

Acknowledgements Open access funding provided by Aalto University. This project has received funding from the European Union’s Horizon 2020 Research and Innovation programme under Grant Agreement No. 671555 and from the Academy of Finland (Finnish Centre of Excellence in Computational Inference Research COIN and Grants 294238, 292334). The authors gratefully acknowledge the computational resources provided by the Aalto Science-IT project and by IT4I under Project ID OPEN-11-20. We also thank Tom Vander Aa for his help in setting up and configuring their distributed parallel BPMF software (Vander Aa et al. 2017) on the Salomon cluster, and the authors of Şimşekli et al. (2017) for sharing their software for NMF + PSGLD. Finally, we wish to thank the four anonymous reviewers for their constructive feedback, which lead to many improvements of the paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

A Proof of Theorem 1

The claim of the theorem is easily verified by rewriting the equation using the factorizations in Eqs. (10)–(12):

$$\begin{aligned}
 p(\mathbf{X}, \mathbf{W} | \mathbf{Y}) &\propto \\
 &p(\mathbf{X}^{(1)}) p(\mathbf{W}^{(1)}) p(\mathbf{Y}^{(1,1)} | \mathbf{X}^{(1)}, \mathbf{W}^{(1)}) \\
 &\times \prod_{i=2}^I \left[p(\mathbf{X}^{(i)}) p(\mathbf{Y}^{(i,1)} | \mathbf{X}^{(i)}, \mathbf{W}^{(1)}) p(\mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)}) p(\mathbf{W}^{(1)} | \mathbf{Y}^{(1,1)})^{-1} \right] \\
 &\times \prod_{j=2}^J \left[p(\mathbf{W}^{(j)}) p(\mathbf{Y}^{(1,j)} | \mathbf{X}^{(1)}, \mathbf{W}^{(j)}) p(\mathbf{X}^{(1)} | \mathbf{Y}^{(1,1)}) p(\mathbf{X}^{(1)} | \mathbf{Y}^{(1,1)})^{-1} \right] \\
 &\times \prod_{i=2}^I \prod_{j=2}^J \left[p(\mathbf{Y}^{(i,j)} | \mathbf{X}^{(i)}, \mathbf{W}^{(j)}) p(\mathbf{X}^{(i)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)}) p(\mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)}) \right. \\
 &\quad \left. \times p(\mathbf{X}^{(i)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(i,1)})^{-1} p(\mathbf{W}^{(j)} | \mathbf{Y}^{(1,1)}, \mathbf{Y}^{(1,j)})^{-1} \right] \\
 &= \prod_{i=1}^I p(\mathbf{X}^{(i)}) \prod_{j=1}^J p(\mathbf{W}^{(j)}) \prod_{i=1}^I \prod_{j=1}^J p(\mathbf{Y}^{(i,j)} | \mathbf{X}^{(i)}, \mathbf{W}^{(j)}).
 \end{aligned}$$

B Hardware platform

The experiments for our method and embarrassingly parallel MCMC were performed on Triton (Aalto Science-IT), a cluster with more than 200 compute nodes. The nodes used for the experiments are equipped with dual 12-core Xeon E5 2680 v3, a clock speed 2.50GHz and 128 GB of RAM. One node was used for posterior inference for each subset. The experiments for distributed/parallel implementations of BMF and NMF were conducted on Salomon (IT4Innovations), a cluster with 576 nodes, each equipped with dual 12-core Intel Xeon E5-2680v3, a clock speed 2.50GHz and at least 128 GB of physical memory per node.

C Wall-clock time speed-up

The results for RMSE versus different partition schemes in Figs. 2 and 3 provide a trade-off between predictive performance and computation time. In this section, we revisit the results for simulated data and MovieLens-1M to analyse the scaling behaviour of our method by evaluating wall-clock time speed-up with respect to the number of workers. We follow Lian et al. (2015) and compute the wall-clock time speed-up (WTS) as follows:

$$\text{WTS} = \frac{\text{Wall-clock time of model on the full data}}{\text{Wall-clock time of distributed method}} \quad (15)$$

The results of this analysis are given in Fig. 6a and Table 3. The theoretical scaling behaviour for an arbitrary fixed problem size is shown in Fig. 6b, where we use the definition of computation time derived in Sect. 4.3 as a function of the number of workers. We conclude that the general behaviour in the empirical results is well in line with what we expect from our discussion in Sect. 4.3, with the actual scale factor depending on the problem-specific configurations (e.g. data size, partition scheme), implementation and hardware.




Fig. 6 Scaling behaviour of our method **a** on simulated data with rows and columns in a decreasing order; **b** theoretical scaling behaviour. The x-axis represents the maximum number of parallel workers that can be used by our method for different partition schemes, the left y-axis is the wall-clock time (solid curve), and the right y-axis is the speed-up (dashed curve)

Table 3 Wall-clock time speed-up on MovieLens-1M and simulated data

Partition	#Workers	Decreasing order			Random order		
		RMSE	Wall-clock time (s)	WTS	RMSE	Wall-clock time (s)	WTS
MovieLens-1M							
30 × 30	841	0.8881	1473	23.056	0.9194	1363	24.419
10 × 10	81	0.8603	4529	7.498	0.8896	3710	8.971
5 × 5	16	0.8512	10,398	3.266	0.8735	8583	3.878
3 × 3	4	0.8485	18,894	1.797	0.8605	16,195	2.055
1 × 1	1	0.8470	33,956	1.0	0.8470	33,283	1.0
Simulated data							
30 × 30	841	1.018	1357	87.069	1.020	1252	94.097
10 × 10	81	1.009	4069	29.029	1.009	4037	29.186
5 × 5	16	1.008	11,542	10.234	1.008	11,467	10.276
3 × 3	4	1.008	28,749	4.109	1.008	27,873	4.227
1 × 1	1	1.008	118,124	1.0	1.008	117,830	1.0

References

- Adams, R., Dahl, G., & Murray, I. (2010). Incorporating side information in probabilistic matrix factorization with Gaussian processes. In *Proceedings of the 26th annual conference on uncertainty in artificial intelligence* (pp. 1–9). AUAI Press.
- Ahn, S., Korattikara, A., Liu, N., Rajan, S., & Welling, M. (2015). Large-scale distributed Bayesian matrix factorization using stochastic gradient MCMC. In *Proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining* (pp. 9–18). <https://doi.org/10.1145/2783258.2783373>.
- Angelino, E., Johnson, M. J., & Adams, R. P. (2016). Patterns of scalable Bayesian inference. *Foundations and Trends in Machine Learning*, 9(2–3), 119–247. <https://doi.org/10.1561/22000000052>.
- Bento, A., Gaulton, A., Hersey, A., Bellis, L., Chambers, J., Davies, M., et al. (2014). The ChEMBL bioactivity database: An update. *Nucleic Acids Research*, 42(D1), D1083–D1090.
- Bhattacharya, A., & Dunson, D. B. (2011). Sparse Bayesian infinite factor models. *Biometrika*, 98(2), 291–306. <https://doi.org/10.1093/biomet/asr013>.
- Cobanoglu, M. C., Liu, C., Hu, F., Oltvai, Z. N., & Bahar, I. (2013). Predicting drug–target interactions using probabilistic matrix factorization. *Journal of Chemical Information and Modeling*, 53(12), 3399–3409. <https://doi.org/10.1021/ci400219z>.
- Comiter, M., Cha, M., Kung, H. T., & Teerapittayanon, S. (2016). Lambda means clustering: Automatic parameter search and distributed computing implementation. In *2016 23rd international conference on pattern recognition (ICPR)* (pp. 2331–2337). <https://doi.org/10.1109/ICPR.2016.7899984>.
- Harper, F. M., & Konstan, J. A. (2015). The MovieLens datasets: History and context. *ACM Transactions on Interactive Intelligent Systems*, 5(4), 19:1–19:18.
- Hernández-Lobato, J. M., Houlisby, N., & Ghahramani, Z. (2014). Probabilistic matrix factorization with non-random missing data. In *Proceedings of the 31st international conference on international conference on machine learning* (Vol. 32, pp. II-1512–II-1520).
- Koren, Y., Bell, R., & Volinsky, C. (2009). Matrix factorization techniques for recommender systems. *Computer*, 42(8), 30–37. <https://doi.org/10.1109/MC.2009.263>.
- Lian, X., Huang, Y., Li, Y., & Liu, J. (2015). Asynchronous parallel stochastic gradient for nonconvex optimization. In *Proceedings of the 28th international conference on neural information processing systems* (Vol. 2, pp. 2737–2745).
- Lopes, H. F., & West, M. (2004). Bayesian model assessment in factor analysis. *Statistica Sinica*, 14(1), 41–67.
- Minsker, S., Srivastava, S., Lin, L., & Dunson, D. (2014). Robust and scalable Bayes via a median of subset posterior measures. [arXiv:1403.2660](https://arxiv.org/abs/1403.2660).
- Murphy, K. P. (2007). *Conjugate Bayesian analysis of the Gaussian distribution*. Technical report, Google

- Neiswanger, W., Wang, C., & Xing, E. (2014). Asymptotically exact, embarrassingly parallel MCMC. In *Proceedings of the 30th conference on uncertainty in artificial intelligence* (pp. 623–632). AUAI Press.
- Nemeth, C., & Sherlock, C. (2017). Merging MCMC subposteriors through Gaussian-process approximations. *Bayesian Analysis*. <https://doi.org/10.1214/17-BA1063> (advance publication).
- Park, S., Kim, Y. D., & Choi, S. (2013). Hierarchical Bayesian matrix factorization with side information. In *Proceedings of the 23rd international joint conference on artificial intelligence* (pp. 1593–1599). AAAI Press.
- Porteous, I., Asuncion, A., & Welling, M. (2010). Bayesian matrix factorization with side information and Dirichlet process mixtures. In *Proceedings of the 24th AAAI conference on artificial intelligence* (pp. 563–568). AAAI Press.
- Salakhutdinov, R., & Mnih, A. (2008a). Bayesian probabilistic matrix factorization using Markov chain Monte Carlo. In *Proceedings of the 25th international conference on machine learning* (pp. 880–887). ACM. <https://doi.org/10.1145/1390156.1390267>.
- Salakhutdinov, R., & Mnih, A. (2008b). Probabilistic matrix factorization. In *Advances in neural information processing systems 20*. MIT Press.
- Scott, S. L., Blocker, A. W., Bonassi, F. V., Chipman, H. A., George, E. I., & McCulloch, R. E. (2016). Bayes and big data: the consensus Monte Carlo algorithm. *International Journal of Management Science and Engineering Management*, 11(2), 78–88. <https://doi.org/10.1080/17509653.2016.1142191>.
- Simm, J., Arany, A., Zakeri, P., Haber, T., Wegner, J. K., Chupakhin, V., et al. (2015). Macau: Scalable Bayesian multi-relational factorization with side information using MCMC. [arXiv:1509.04610](https://arxiv.org/abs/1509.04610).
- Şimşekli, U., Durmus, A., Badeau, R., Richard, G., Moulines, É., & Cemgil, A. T. (2017). Parallelized stochastic gradient Markov chain Monte Carlo algorithms for non-negative matrix factorization. In *IEEE International Conference on Acoustics, Speech and Signal Processing* (pp. 2242–2246). ICASSP 2017.
- Şimşekli, U., Koptagel, H., GÜltaş, H., Cemgil, A. T., Öztoprak, F., & Birbil, Ş. I. (2015). Parallel stochastic gradient MCMC for matrix factorisation models. [arXiv:1506.01418](https://arxiv.org/abs/1506.01418)
- Srivastava, S., Li, C., & Dunson, D. (2015). Scalable Bayes via barycenter in Wasserstein space. [arXiv:1508.05880](https://arxiv.org/abs/1508.05880).
- Teh, Y. W. (2007). *Exponential Families: Gaussian, Gaussian-Gamma, Gaussian-Wishart*. Multinomial, Technical report, University College London.
- Vander Aa, T., Chakroun, I., & Haber, T. (2016). Distributed Bayesian probabilistic matrix factorization. In *2016 IEEE international conference on cluster computing* (pp. 346–349). IEEE Computer Society.
- Vander Aa, T., Chakroun, I., & Haber, T. (2017). Distributed Bayesian probabilistic matrix factorization. *Procedia Computer Science*, 108, 1030–1039.
- Vehtari, A., Gelman, A., Sivula, T., Jylänki, P., Tran, D., Sahai, S. et al. (2018). Expectation propagation as a way of life: A framework for Bayesian inference on partitioned data. [arXiv:1412.4869](https://arxiv.org/abs/1412.4869).
- Wang, X., Dunson, D. (2013). Parallelizing MCMC via Weierstrass sampler. [arXiv:1312.4605](https://arxiv.org/abs/1312.4605).
- Wang, X., Guo, F., Heller, K. A., Dunson, D. B. (2015). Parallelizing MCMC with random partition trees. In *Advances in neural information processing systems 28* (pp. 451–459). Curran Associates, Inc.
- Wang, X., Peng, P., Dunson, D. B. (2014). Median selection subset aggregation for parallel inference. In *Advances in neural information processing systems 27* (pp. 2195–2203). Curran Associates, Inc.
- Xu, M., Lakshminarayanan, B., Teh, Y. W., Zhu, J., & Zhang, B. (2014). Distributed Bayesian posterior sampling via moment sharing. In Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N. D., Weinberger, K. Q. (Eds.), *Advances in Neural Information Processing Systems 27* (pp. 3356–3364). Curran Associates, Inc.. <http://papers.nips.cc/paper/5596-distributed-bayesian-posterior-sampling-via-moment-sharing.pdf>.