Large-Scale Distributed Bayesian Matrix Factorization using Stochastic Gradient MCMC

Ahn, S.; Korattikara, A.; Liu, N.; Rajan, S.; Welling, M.

Link to publication

Large-Scale Distributed Bayesian Matrix Factorization using Stochastic Gradient MCMC

Sungjin Ahn
University of California, Irvine
sungjia@ics.uci.edu

Anoop Korattikara [*]
Google
kbanoop@google.com

Nathan Liu
Yahoo Labs
nanliu@yahoo-inc.com

Suju Rajan
Yahoo Labs
suju@yahoo-inc.com

Max Welling [†]
University of Amsterdam
m.welling@uva.nl

ABSTRACT

Despite having various attractive qualities such as high prediction accuracy and the ability to quantify uncertainty and avoid overfitting, Bayesian Matrix Factorization has not been widely adopted because of the prohibitive cost of inference. In this paper, we propose a scalable distributed Bayesian matrix factorization algorithm using stochastic gradient MCMC. Our algorithm, based on Distributed Stochastic Gradient Langevin Dynamics, can not only match the prediction accuracy of standard MCMC methods like Gibbs sampling, but at the same time is as fast and simple as stochastic gradient descent. In our experiments, we show that our algorithm can achieve the same level of prediction accuracy as Gibbs sampling an order of magnitude faster. We also show that our method reduces the prediction error as fast as distributed stochastic gradient descent, achieving a 4.1% improvement in RMSE for the Netflix dataset and an 1.8% for the Yahoo music dataset.

1. INTRODUCTION

Recommender systems have become a pervasive tool in industry to understand customers and their interests in products. Examples range between music recommendation (Pandora), book recommendation (Amazon), movie recommendation (Netflix), news recommendation (Yahoo) to partner recommendation (eHarmony). Recommender systems represent a personalized technology that can help filter at an individual level the enormous amounts of information that is available to us. Given the exponential growth of data, recommender systems are likely to play an increasingly important role to manage our information streams.

During 2006-2011 Netflix [2, 7] ran a competition where teams around the world could develop and test new recommender technology on Netflix movie rating data. A few valuable lessons were

---

learnt from that exercise. First, matrix factorization methods work very well compared to nearest neighbor type models. Second, averaging over many different models pays off in terms of prediction accuracy. One particularly effective model was Bayesian probabilistic matrix factorization (BPMF) [27] where predictions are averaged over samples from the posterior distribution. Besides improved prediction accuracy, a full Bayesian analysis also comes with additional advantages such as probabilities over models, confidence intervals, robustness against overfitting, and incorporating prior knowledge and side-information [4, 24].

Unfortunately, since the number of user-product interactions can easily run into the billions, posterior inference is usually too expensive to be practical. Learning at that scale requires data and computation to be distributed over many machines and learning updates to only depend on small minibatches of the data. Effective distributed learning algorithms have been devised for alternating least squares (ALS) and stochastic gradient descent (SGD) [13, 25, 30, 28, 22, 15, 19, 18, 31]. In particular, Distributed Stochastic Gradient Descent (DSGD) [13] has achieved a significant speed-up by assigning partitioned rating matrix blocks to workers and then by updating some "orthogonal" blocks in parallel using "stratified" SGD. DSGD outperformed other parallel SGD approaches such as PSGD [15, 19] and ISGD [18, 31] where SGD is applied also on some subsets of the ratings while synchronizing globally after each sub-epoch (PSGD) or once at the end of the training (ISGD). Unfortunately, so far it has proven difficult to apply these advances in distributed learning to posterior sampling in Bayesian matrix factorization models. For instance, for BPMF which requires $\mathcal{O}((L + M)D^3)$ computation per iteration (with L and M are number of users and items, and D is latent feature dimension), distributed computation has not nearly been as effective.

In this paper, we propose a scalable and distributed Bayesian matrix factorization method which combines the predictive accuracy of Bayesian inference and the learning efficiency of stochastic gradient updates. To this end, we extend a recently developed MCMC method, called Stochastic Gradient Langevin Dynamics (SGLD) [29], so that the updates become efficient in the setting of distributed, large-scale matrix factorization. We adapt the SGLD updates to make them suitable for distributed learning on subsets of users and products (or blocks). Each worker manages only a small block of the rating matrix, and updates and communicates only a small subset of the parameters in a fully-asynchronous or weakly-synchronous fashion. Unlike distributed SGD where a single model is learnt, our method deploys multiple parallel chains over workers. Consequently, samples are collected at a much faster rate than ordi-

nary MCMC and the multiple parallel chains can explore different modes of parameter space. Both features are reducing the variance and increasing the accuracies of our predictions.

In the experiments on the Netflix and Yahoo music datasets (the latter being one of the largest publicly available dataset for recommendation problems), we show that our method achieves the same level of accuracy as BPMF but an order of magnitude faster. Reversely, at almost the same efficiency as distributed SGD, our method achieves much better accuracy (4.1% RMSE improvement for the Netflix dataset and 1.8% for Yahoo music dataset). As such we believe that the method proposed in this paper is currently the most competitive matrix factorization method for industry scale problems.

2. PRELIMINARIES

2.1 Bayesian Matrix Factorization

Suppose we have L users and M items. Our goal is to learn latent feature vectors $U_i, V_j \in \mathbb{R}^D$ such that the rating R_{ij} for item j by user i can be predicted as $R_{ij} \approx U_i^\top V_j$. We denote the entire rating matrix by $\mathbf{R} \in \mathbb{R}^{L \times M}$, and the latent feature matrices by $\mathbf{U} \in \mathbb{R}^{D \times L}$ and $\mathbf{V} \in \mathbb{R}^{D \times M}$, so that $\mathbf{R} \approx \mathbf{U}^\top \mathbf{V}$. Assuming a Gaussian error model, the likelihood of the parameters $\mathbf{U}$ and $\mathbf{V}$ can be written as:

$$p(\mathbf{R}|\mathbf{U}, \mathbf{V}, \tau) = \prod_{i=1}^{L} \prod_{j=1}^{M} \left[\mathcal{N}(R_{ij}|U_i^\top V_j, \tau^{-1}) \right]^{I_{ij}}. \quad (1)$$

where I_{ij} is equal to 1 if user i rated item j and 0 otherwise. Throughout the paper, we fixed $\tau = 1$ for simplicity[1]. Although, in theory, $\mathbf{U}$ and $\mathbf{V}$ can be learned by maximizing the likelihood above, this results in severe over-fitting because only a few ratings are known (i.e. $\mathbf{R}$ is very sparse).

Therefore, a Bayesian Probabilistic Matrix Factorization (BPMF) model was proposed to overcome this problem [27]. In addition to controlling over-fitting through posterior averaging, BPMF also provides estimates of uncertainty through the posterior predictive distribution. The BPMF model as proposed in [27] is as follows. We place priors on $\mathbf{U}$ and $\mathbf{V}$ as:

$$p(\mathbf{U}|\mu_U, \Lambda_U) = \prod_{i=1}^{L} \mathcal{N}(U_i|\mu_U, \Lambda_U^{-1}), \quad (2)$$

$$p(\mathbf{V}|\mu_V, \Lambda_V) = \prod_{j=1}^{M} \mathcal{N}(V_j|\mu_V, \Lambda_V^{-1}). \quad (3)$$

We further place Gaussian-Wishart hyper-priors on the user and item hyperparameters $\Theta_U = \{\mu_U, \Lambda_U\}$ and $\Theta_V = \{\mu_V, \Lambda_V\}$:

$$p(\Theta_U|\Theta_0) = \mathcal{N}(\mu_U|\mu_0, (\beta_0\Lambda_U)^{-1})\mathcal{W}(\Lambda_U|W_0, \nu_0), \quad (4)$$

$$p(\Theta_V|\Theta_0) = \mathcal{N}(\mu_V|\mu_0, (\beta_0\Lambda_V)^{-1})\mathcal{W}(\Lambda_V|W_0, \nu_0), \quad (5)$$

where ν_0 is the number of degrees of freedom and W_0 is a $D \times D$ scale matrix. We collectively denote the parameters of the hyperprior by $\Theta_0 = \{\mu_0, \beta_0, \nu_0, W_0\}$.

At test time, the predictive distribution of an unknown rating R_{ij}^* can be obtained by marginalizing over both model parameters $\mathbf{U}, \mathbf{V}$ and hyper-parameters Θ_U, Θ_V,

$$p(R_{ij}^*|\mathbf{R}, \Theta_0) = \int\int p(R_{ij}^*|U_i, V_j)p(\mathbf{U}, \mathbf{V}|\mathbf{R}, \Theta_U, \Theta_V)$$
$$p(\Theta_U, \Theta_V|\Theta_0)\mathrm{d}\{\mathbf{U}, \mathbf{V}\}\mathrm{d}\{\Theta_U, \Theta_V\} \quad (6)$$

[1]All update equations are derived with $\tau = 1$.

Algorithm 1 Gibbs Sampling for BPMF
1: Initialize model parameters $\mathbf{U}^{(1)}, \mathbf{V}^{(1)}$
2: **for** $t = 1 : T$ **do**
3:   // Sample hyperparameters
   $\Theta_U^{(t)} \sim p(\Theta_U|\mathbf{U}^{(t)}, \Theta_0), \quad \Theta_V^{(t)} \sim p(\Theta_V|\mathbf{V}^{(t)}, \Theta_0)$
4:   **for** $i = 1 : L$ **in parallel do**
5:     $U_i^{(t+1)} \sim p(U_i|\mathbf{R}, \mathbf{V}^{(t)}, \Theta_U^{(t)})$ // sample user features
6:   **end for**
7:   **for** $j = 1 : M$ **in parallel do**
8:     $V_j^{(t+1)} \sim p(V_j|\mathbf{R}, \mathbf{U}^{(t)}, \Theta_V^{(t)})$ // sample item features
9:   **end for**
10: **end for**

We can estimate this using a Monte Carlo approximation:

$$p(R_{ij}^*|\mathbf{R}, \Theta_0) \approx \frac{1}{T} \sum_{t=1}^{T} p(R_{ij}^*|U_i^{(t)}, V_j^{(t)}). \quad (7)$$

where $\left\{U_i^{(t)}, V_j^{(t)}\right\}$ is the t-th sample from the posterior distribution:

$$p(\mathbf{U}, \mathbf{V}, \Theta_U, \Theta_V|\mathbf{R}, \Theta_0). \quad (8)$$

These samples can be generated using Gibbs sampling (Algorithm 1), since by conjugacy the conditional distributions of U_i and V_j are Gaussian, and those of Θ_U and Θ_V are Gaussian-Wishart. However, sampling from the conditional distribution of U_i or V_j involves $\mathcal{O}(D^3)$ computations (for inverting a $D \times D$ precision matrix) and since this has to be done for each user and item, results in a total of $\mathcal{O}((L+M)D^3)$ computations per iteration. Thus, BPMF using Gibbs sampling cannot scale up to real world recommender systems with millions of users and / or items.

Although it is possible to parallelize BPMF using MapReduce style synchronous global updates, the cubic order complexity still limits its applicability to small D. Also, we require a large number of workers to effectively distribute the $L + M$ cubic-order computations. Furthermore, since running the Gibbs sampler from scratch is too expensive, a separate SGD optimizer is usually deployed to reach near the Maximum-a-Posteriori (MAP) state before starting the Gibbs sampler. However, running two different large-scale distributed algorithms, each of which requires different optimal settings for the distribution of data and parameters, as well as cluster architectures, adds another considerable level of complexity.

2.2 Stochastic Gradient Langevin Dynamics

Assume we have a dataset of N i.i.d. data points, denoted by $\mathcal{X} = \{x_n\}_{n=1}^N$, which we model using a distribution $p(x|\theta)$ parameterized by $\theta \in \mathbb{R}^D$. We choose a prior distribution $p(\theta)$ and our goal is to sample from the posterior distribution $p(\theta|\mathcal{X}) \propto p(\mathcal{X}|\theta)p(\theta)$ using MCMC.

One way of obtaining efficient MCMC proposals is to use the gradient of the target density [26, 12, 21, 14], e.g. Langevin Dynamics [26] is an MCMC algorithm which proposes candidate states according to:

$$\theta_{t+1} \leftarrow \theta_t + \frac{\epsilon_t}{2} \left\{ \nabla_{\theta_t} \log p(\theta_t) + \sum_{x \in \mathcal{X}} g(\theta; x) \right\} + \nu_t$$
$$\text{where } \nu_t \sim \mathcal{N}0, \epsilon_t I \quad (9)$$

In the above, ϵ_t is the step size and $g(\theta; x) = \nabla_\theta \log p(x|\theta)$ is the score. A Metropolis-Hastings (MH) test is then used to decide whether to accept or reject the proposal. The gradient information allows the Langevin algorithm to make proposals to high density

regions and therefore have a high probability of acceptance. However, in large-scale problems where $N = |\mathcal{X}|$ can be very large, the $\mathcal{O}(N)$ computations per update, required for computing the gradient as well as for the MH test, is infeasible.

Stochastic Gradient Langevin Dynamics (SGLD) [29] is the first in a line of recently developed approximate MCMC algorithms [5, 23, 9, 10, 6] that try to address this issue using noisy gradients that can be cheaply computed from a mini-batch of $n \ll N$ data points. SGLD uses the following update rule:

$$\theta_{t+1} \leftarrow \theta_t + \frac{\epsilon_t}{2} \left\{ \nabla_{\theta_t} \log p(\theta_t) + N\bar{g}(\theta_t; \mathcal{M}_t) \right\} + \nu_t. \quad (10)$$

Here $\bar{g}(\theta_t; \mathcal{M}_t) = \frac{1}{n} \sum_{x \in \mathcal{M}_t} g(x; \theta_t)$, the mean score computed from a mini-batch $\mathcal{M}_t$. SGLD converges to the true posterior distribution if the step size is annealed to zero at a rate that satisfies the following conditions:

$$\sum_{t=1}^{\infty} \epsilon_t = \infty, \qquad \sum_{t=1}^{\infty} \epsilon_t^2 < \infty. \quad (11)$$

SGLD does not use accept-reject tests because the acceptance rate tends to one as the step size goes to zero. Therefore, unlike traditional MCMC algorithms which require $\mathcal{O}(N)$ computations per iteration, SGLD requires only $\mathcal{O}(n)$ computations.

More generally, it is valid to replace $\bar{g}(\theta_t; \mathcal{M}_t)$ in eqn. 10 with any estimator $f(\theta, Z; X)$ that satisfies the following conditions: (i) it is an unbiased estimator of the true gradient i.e. $\mathbb{E}_Z[f(\theta, Z; \mathcal{X})] = \bar{g}(\theta; \mathcal{X})$ (ii) it has finite variance $\mathbb{V}_Z[f(\theta, Z; \mathcal{X})] < \infty$. Here, the expectation and variance are w.r.t. the distribution $p(Z; \mathcal{X})$ of the auxiliary random variable Z.

Distributed SGLD (DSGLD) [6] further extends the power of stochastic gradient MCMC using distributed computing. In DS-GLD, the dataset is first partitioned and distributed to S workers. Then, multiple chains collect samples in parallel by sampling for the length of a round (called a *trajectory*) at a worker. After a round, each chain switches to a different worker. In [6], it is shown that using the following valid SGLD update rule, we can collect samples from the posterior using the distributed datasets:

$$\theta_{t+1} \leftarrow \theta_t + \frac{\epsilon_t}{2} \left\{ \nabla_{\theta_t} \log p(\theta_t) + \frac{N^{(s)}}{v^{(s)}} \bar{g}(\theta_t; \mathcal{M}_t^{(s)}) \right\} + \nu_t. \quad (12)$$

Here, s is the index of the worker where a chain resides at iteration t, $N^{(s)}$ is the size of the local dataset at worker s, and $v^{(s)}$ is the normalized visiting rate to worker s such that $\sum_s \nu^{(s)} = 1$ and $\nu^{(s)} \in (0, 1)$. The mini-batch $\mathcal{M}_t^{(s)}$ is sampled only from the local dataset of worker s.

3. BAYESIAN MATRIX FACTORIZATION USING SGLD

3.1 Model

We will now show how DSGLD can be used for BPMF. Instead of the model described in Section 2.1, we will use a slightly simplified model [20, 9]. We use the same likelihood as in eqn. 1, but choose simpler priors:

$$p(\mathbf{U}|\Lambda_U) = \prod_{i=1}^{L} \mathcal{N}(U_i|0, \Lambda_U^{-1}), \quad (13)$$

$$p(\mathbf{V}|\Lambda_V) = \prod_{j=1}^{M} \mathcal{N}(V_j|0, \Lambda_V^{-1}). \quad (14)$$

Here, Λ_U and Λ_V are D-dimension diagonal matrices whose d-th diagonal elements are λ_{U_d} and λ_{V_d} respectively. We also choose the following hyper-priors:

$$\lambda_{U_d}, \lambda_{V_d} \sim \text{Gamma}(\alpha_0, \beta_0). \quad (15)$$

We choose this simplified model because the proposed method benefits mainly from performing a large number of inexpensive updates (i.e. collecting many samples) per unit time rather than very expensive but high quality updates. The above model is well suited for this because each latent vector can be updated in linear ($\mathcal{O}(D)$) time. At the same time, we still benefit from the power of Bayesian inference through marginalization of the important regularization parameters $\Lambda = \{\Lambda_U, \Lambda_V\}$ as well as $\mathbf{U}$ and $\mathbf{V}$.

Although it is possible to apply our method to the model in Section 2.1, updating the full covariance matrix is more expensive ($\mathcal{O}(D^2)$ time per update) and therefore requires more time to converge without significant gain in accuracy (as per our pilot experiments).

3.2 Inference

In the following section, we first present our algorithm in a single machine setting and later extend it for distributed inference. We alternate between sampling from $p(\mathbf{U}, \mathbf{V}|\mathbf{R}, \Lambda)$ using SGLD and sampling from $p(\Lambda|\mathbf{R}, \mathbf{U}, \mathbf{V})$ using Gibbs.

3.2.1 Sampling $\mathbf{U}$, $\mathbf{V}$ | Λ, $\mathbf{R}$ using SGLD

Since, usually only $N \ll M \times L$ ratings are observed, the rating matrix $\mathbf{R}$ is stored using a sparse representation as $\mathcal{X} = \{x_n = (p_n, q_n, r_n)\}_{n=1}^{N}$, where each x_n is a (user, item, rating) tuple and N is the number of observed ratings. The gradient of the log-posterior w.r.t.[2] U_i is:

$$G(\mathcal{X}) = \sum_{n=1}^{N} g_n(U_i; \mathcal{X}) - \Lambda_U U_i \quad (16)$$

where

$$g_n(U_i; \mathcal{X}) = \mathbb{I}[p_n = i|\mathcal{X}](r_n - U_{p_n}^\top V_{q_n})V_{q_n} \quad (17)$$

Here $\mathbb{I}[p_n = i|\mathcal{X}]$ is an indicator function that equals 1 if the n-th tuple in $\mathcal{X}$ pertains to user i and 0 otherwise. To use SGLD, we need an unbiased estimate of this gradient that can be computed cheaply from a mini-batch.

One way to obtain this is by subsampling a mini-batch $\mathcal{M} = \{(p_n, q_n, r_n)\}_{n=1}^{m}$ of m tuples from $\mathcal{X}$ and computing the following stochastic approximation of the gradient:

$$G_1(\mathcal{M}) = N\bar{g}(U_i; \mathcal{M}) - \Lambda_U U_i \quad (18)$$

where, $\bar{g}(U_i; \mathcal{M}) = \frac{1}{m} \sum_{n=1}^{m} g_n(U_i; \mathcal{M})$. Note that the mini-batch is subsampled from the complete dataset $\mathcal{X}$ and not just from the tuples associated with user i. The expectation of G_1 over all possible mini-batches is:

$$\begin{aligned} \mathbb{E}_\mathcal{M}[G_1(\mathcal{M})] &= \mathbb{E}_\mathcal{M}\left[N\bar{g}(U_i; \mathcal{M})\right] - \Lambda_U U_i \\ &= \sum_{n=1}^{N} g_n(U_i; \mathcal{X}) - \Lambda_U U_i \\ &= G(\mathcal{X}). \end{aligned}$$

Since G_1 is an unbiased estimator of the true gradient, we can use it for computing SGLD updates. However, note that G_1 is non-zero even for users that are not in the mini-batch $\mathcal{M}$, because of

---

[2]We derive only w.r.t. U_i. Update rules for other parameters can be obtained by the same procedure.

1	2
3	4

(a) square

1
2
3
4

(b) column

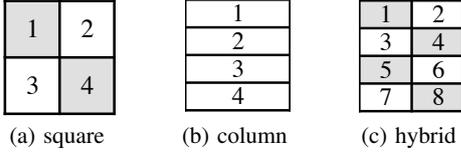| 1	2
3	4
5	6
7	8

(c) hybrid

Figure 1: Block split schemes.

the prior gradient term $-\Lambda_U U_i$. Therefore, we have to update the parameters for all users in every iteration, which is very expensive.

If we were to update only the parameters of users who have ratings in the mini-batch $\mathcal{M}$, the estimator can be written as:

$$G_2(\mathcal{M}) = N\bar{g}(U_i; \mathcal{M}) - \mathbb{I}[i \in \mathcal{M}_p]\Lambda_U U_i \qquad (19)$$

where $\mathbb{I}[i \in \mathcal{M}_p]$ is equal to 1 if $\mathcal{M}$ contains a tuple associated with user i and 0 otherwise. However, G_2 is not an unbiased estimator of the true gradient:

$$\mathbb{E}_{\mathcal{M}}[G_2(\mathcal{M})] = \sum_{n=1}^{N} g_n(U_i; \mathcal{X}) - h_{i*}\Lambda_U U_i. \qquad (20)$$

where $h_{i*} = \mathbb{E}_{\mathcal{M}}[\mathbb{I}[i \in \mathcal{M}_p]]$, i.e. the fraction of mini-batches that contains at least one tuple associated with user i (among all possible mini-batches). If the mini-batches are sampled with replacement, we can compute this as:

$$h_{i*} = 1 - \left(1 - \frac{N_{i*}}{N}\right)^m \qquad (21)$$

where $N_{i*} = \sum_{n=1}^{N} \mathbb{I}[p_n = i | \mathcal{X}]$, the number of ratings by user i in the complete dataset $\mathcal{X}$. Thus, we can remove the bias in G_2 by multiplying the gradient of the prior term with h_{i*}^{-1} as follows:

$$G_3(\mathcal{M}) = N\bar{g}(U_i; \mathcal{M}) - \mathbb{I}[i \in \mathcal{M}_p]h_{i*}^{-1}\Lambda_U U_i. \qquad (22)$$

G_3 is an unbiased estimator of the true gradient G and is non-zero only for users that have at least one rating in $\mathcal{M}$. Thus we need to update only a subset of user features in each iteration. The SGLD update rule (for users with ratings in $\mathcal{M}_t$) is:

$$U_{i,t+1} \leftarrow U_{i,t} + \frac{\epsilon_t}{2}\left\{N\bar{g}(U_{i,t}; \mathcal{M}_t) - \frac{\Lambda_U U_{i,t}}{h_{i*}}\right\} + \nu_t \qquad (23)$$

3.2.2 Sampling $\Lambda | \mathbf{U}, \mathbf{V}$

We can easily sample from the conditional $p(\Lambda | \mathbf{U}, \mathbf{V})$, because by conjugacy:

$$\lambda_{U_d} | \mathbf{U}, \mathbf{V} \sim \text{Gamma}\left(\alpha_0 + \frac{L}{2}, \beta_0 + \frac{1}{2}\sum_{i=1}^{L} U_{di}^2\right), \qquad (24)$$

$$\lambda_{V_d} | \mathbf{U}, \mathbf{V} \sim \text{Gamma}\left(\alpha_0 + \frac{M}{2}, \beta_0 + \frac{1}{2}\sum_{i=1}^{M} V_{dj}^2\right). \qquad (25)$$

If this is computationally demanding, we can also consider updating Λ using SGLD or mini-batch Metropolis-Hastings [16].

3.3 Distributed Inference

For distributed inference, we partition the rating matrix $\mathbf{R}$ into a number of blocks. Fig. 1 shows a few different ways of partitioning $\mathbf{R}$. Two blocks are said to be *orthogonal* to each other if the users and items in one block do not appear in the other block. A set of two or more mutually orthogonal blocks is called an *orthogonal block group* (or simply, orthogonal group). For example, the two

gray-colored blocks (1 and 4) in Fig. 1 (a) are orthogonal to each other and thus form an orthogonal group. In Fig. 1 (b), the blocks are not orthogonal because all columns are shared. In this case, we say that each block by itself is an orthogonal group.

The blocks are then distributed to workers in such a way that all blocks are assigned and a worker has at least one block. In the following, we assume for simplicity that each worker is a single-core machine. However, it is easy to generalize our algorithm to take advantage of multi-core (or threads) workers with shared memory support.

We will now describe our distributed algorithm for BPMF. First, imagine that there is only one Markov chain c (but the dataset is distributed across multiple workers). A central parameter server holds the *global* parameters $\mathbf{U}^c$ and $\mathbf{V}^c$ of chain c. Since Λ depends only on $\mathbf{U}^c$ and $\mathbf{V}^c$, it is easy to update Λ at the parameter server using Gibbs as per Eqns. 25 and 24. Thus, we will focus on the DSGLD part of the chain that samples from $p(\mathbf{U}, \mathbf{V} | \mathbf{R}, \Lambda)$.

Each sampling round consists of the following steps: (1) The parameter server picks a block s via a block-scheduler and sends the corresponding sub-parameter $\mathbf{U}^{(c,s)}$ and $\mathbf{V}^{(c,s)}$ to the block's worker. (2) The worker updates the sub-parameter by running DS-GLD (see section 3.3.1 for update equations) for a number of iterations using its local block of ratings. (3) The worker sends the final sub-parameter state back to the parameter server. (4) The parameter server updates its global copy to the new sub-parameter state.

Thus, the Markov chain jumps among the distributed blocks through the corresponding workers and updates the sub-parameters associated with the block chosen in each round. Since each iteration of local DSGLD updates requires only a mini-batch of data, sampling is very fast. Also, communication overhead is low because a) the multiple local updates (iterations) performed within a round do not require any communication b) only a small sub-parameter associated with a specific block is transferred in each round. There are two levels of parallelization that we use to further speed up sampling.

1. Parallel updates within a chain: . A chain can update sub-parameters $\mathbf{U}^{(c,s_1)}$ and $\mathbf{U}^{(c,s_2)}$ in parallel if the blocks s_1 and s_2 are orthogonal to each other. For example, in Fig. 1 (a), updating block 1 *and then* block 4 produces the same result as updating both in parallel. This makes the algorithm progress faster in terms of number of updated parameters per round. The actual performance improvement is dependent on the size of the orthogonal group. For instance, with a 4×4 split, the algorithm will update the parameters faster than with a 2×2 split because more parameter blocks can be updated in parallel. However, updates in smaller blocks can be noisier, because the gradients computed from smaller blocks will have higher variance. Therefore, at some point the loss in performance caused by noisier updates on small blocks can exceed the gain obtained by faster updating of the parameters.

2. Multiple parallel chains: We can run as many chains in parallel as we like, subject to only computational resource constraints. Each chain can update its parameters in parallel independent of other chains. Hence, the chains are asynchronous in the sense that the status of a chain does not block other chains unless the chains conflict for computation resources. For the split in Fig. 1 (a), one chain can update using the gray block group while another chain is using the white block group. Or both chains can use the same block if we assume a shared memory multi-threaded implementation. By running multiple chains in parallel, we effectively multiply the number of collected samples by the number of parallel chains. Since the variance of an MCMC estimator is inversely proportional to the number of samples, fast sample generation will compensate for the low mixing rate of SGLD. Also, by initializing the different
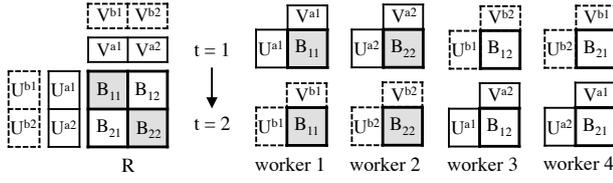
Figure 2: An example illustration. On the left, a matrix $\mathbf{R}$ is partitioned into 2×2 blocks, $\mathbf{B}_{11}, \cdots, \mathbf{B}_{22}$. There are two orthogonal groups (the gray $(\mathbf{B}_{11}, \mathbf{B}_{22})$ group and the white $(\mathbf{B}_{12}, \mathbf{B}_{21})$ group). We run two independent chains, chain a with parameters $\mathbf{U}^a$ and $\mathbf{V}^a$ (solid-line rectangles) and chain b, with parameters $\mathbf{U}^b$ and $\mathbf{V}^b$ (dotted-line rectangles). Given four workers, we assign a block to each worker. At round $t = 1$, chain a updates using the gray orthogonal group and chain b using the white orthogonal group. Note that the entire $\mathbf{U}$ and $\mathbf{V}$ matrices of both chains are updated in this single round. In the next round, the chains are assigned to the next orthogonal groups by the block-scheduler.

Algorithm 2 DSGLD at parameter server

1: Initialize model parameters of each chain $\{\mathbf{U}_1^c, \mathbf{V}_1^c, \Lambda_1^c\}_{c=1}^C$, step sizes $\{\epsilon_t\}$
2: **for** each chain c **parallel do**
3:     **for** t=1:max_iter **do**
4:         $\mathcal{B}_c \leftarrow$ GET_ORTHO_BLOCK_GROUP(c, t)
5:         **for** worker $s \in \mathcal{B}_c$ **do**
6:             $\mathbf{U}_{t+1}^{(c,s)}, \mathbf{V}_{t+1}^{(c,s)} \leftarrow$ WKR_ROUND$(\mathbf{U}_t^{(c,s)}, \mathbf{V}_t^{(c,s)}, \Lambda_t^{(c)}, \epsilon_t)$
7:         **end for**
8:         **if** not burn-in **then**
9:             Store $\mathbf{U}_{t+1}^{(c)}, \mathbf{V}_{t+1}^{(c)}$ as a sample of chain c
10:           Sample $\Lambda_{t+1}^{(c)} | \mathbf{U}_{t+1}^{(c)}, \mathbf{V}_{t+1}^{(c)}$ using Eqn. (24) and (25)
11:         **end if**
12:     **end for**
13: **end for**

---

chains in different places of parameter space, we can explore multiple local minima. This is especially important for large-scale high dimensional problems where the time budget is usually not enough for a single chain to mix between different local minima.

An illustration of these ideas is given in Fig. 2. Algorithms 2 and 3 describe the operations at the parameter server and workers respectively.

A proper block splitting scheme can be chosen according to the characteristics of the problem and available resources. In other words, we can trade-off *within-chain* parallelization and *between-chain* parallelization. For example, given S workers, by using a squared split as in Fig. 1 (a), we can run $\sqrt{S}$ chains in parallel where each chain updates $\sqrt{S}$ blocks in parallel. This way we maximize the within-chain parallelism. On the other hand, by reducing the size of orthogonal groups, we can decrease the within-chain parallelism in order to increase the between-chain parallelization, i.e. number of parallel chains. At an extreme of this approach, we can let each block become an orthogonal group by itself as in Fig. 1 (b) and run S independent chains in parallel. Note that in this case, we can choose not only the column splitting but any splitting scheme. Our experiment results suggest to maximize the within-chain parallelism as the dataset size increases. For smaller datasets, we may benefit more from the generalization performance of a large number of parallel chains than from a smaller number of chains using the block orthogonality.

Algorithm 3 DSGLD at worker s

1: Initialize $\bar{h}_{i*}, \bar{h}_{*j}$, round length γ, mini-batch size m
2: **function** WKR_ROUND$(\mathbf{U}^{(c,s)}, \mathbf{V}^{(c,s)}, \Lambda^{(c)}, \epsilon_t)$
3:     **for** $t = 1 : \gamma$ **do**
4:         Sample a mini-batch $\mathcal{M}_t$ from $\mathcal{X}^{(s)}$
5:         **for** each user i and item j in $\mathcal{M}_t$ **parallel do**
6:             Update U_i, V_j using Eqn. (29) and (30)
7:         **end for**
8:     **end for**
9:     Send updated $\mathbf{U}^{(c,s)}$ and $\mathbf{V}^{(c,s)}$ to the parameter server
10: **end function**

3.3.1 Distributed SGLD

Since $\mathcal{X}$ (the sparse representation of $\mathbf{R}$) is partitioned into S blocks $\mathcal{X}^{(1)}, \ldots, \mathcal{X}^{(S)}$, each worker uses only one of the $\mathcal{X}^{(s)}$ for computing updates. Thus, we need to modify the bias correctors in Eqn. (21) so that the gradient estimator remains unbiased under this constraint. If we assume $\cup_{s=1}^S \mathcal{X}^{(s)} = \mathcal{X}$ and $\cap_{s=1}^S \mathcal{X}^{(s)} = \emptyset$, and that worker s is visited with normalized frequency $v^{(s)}$, the correction factors for users and items can be shown to be, respectively:

$$\bar{h}_{i*} = \sum_{s=1}^S v^{(s)} h_{i*}^{(s)}, \qquad \bar{h}_{*j} = \sum_{s=1}^S v^{(s)} h_{*j}^{(s)} \qquad (26)$$

where:

$$h_{i*}^{(s)} = 1 - \left(1 - \frac{N_{i*}^{(s)}}{N^{(s)}}\right)^m, \quad h_{*j}^{(s)} = 1 - \left(1 - \frac{N_{*j}^{(s)}}{N^{(s)}}\right)^m \qquad (27)$$

here $N^{(s)} = |\mathcal{X}^{(s)}|$, the total number of ratings in s, and

$$N_{i*}^{(s)} = \sum_{n=1}^{N^{(s)}} \mathbb{I}[p_n = i | \mathcal{X}^{(s)}], \quad N_{*j}^{(s)} = \sum_{n=1}^{N^{(s)}} \mathbb{I}[q_n = j | \mathcal{X}^{(s)}]. \qquad (28)$$

i.e. the number of ratings by user i and of item j respectively in s. Therefore, the local DSGLD update rule using block $\mathcal{X}^{(s)}$ is:

$$U_{i,t+1} \leftarrow U_{i,t} + \frac{\epsilon_t}{2} \left\{ \frac{N^{(s)}}{v^{(s)}} \bar{g}(U_{i,t}; \mathcal{M}_t^{(s)}) - \frac{\Lambda_U U_{i,t}}{\bar{h}_{i*}} \right\} + \nu_t \quad (29)$$

$$V_{j,t+1} \leftarrow V_{j,t} + \frac{\epsilon_t}{2} \left\{ \frac{N^{(s)}}{v^{(s)}} \bar{g}(V_{j,t}; \mathcal{M}_t^{(s)}) - \frac{\Lambda_V V_{j,t}}{\bar{h}_{*j}} \right\} + \nu_t. \quad (30)$$

The above rule updates only the sub-parameter associated with block s using only rating tuples in s.

4. EXPERIMENTS

4.1 Algorithms and Models

	Optimization	MCMC
Single Machine	SGD	SGLD, Gibbs
Distributed	DSGD	DSGLD

Table 1: Algorithms.

We compared five algorithms: SGD, DSGD, SGLD, DSGLD, and Gibbs sampling. As shown in Table 1, each algorithm can be classified based on whether it is running on a single machine or a distributed architecture, and also based on whether it is an optimization or MCMC algorithm. Since Gibbs sampling was very slow, we update user/item features in parallel (as suggested in [27]) using multiple cores of a single machine. Thus, by Gibbs sampling

we will mean the parallelized (but not distributed) version from now on.

For DSGLD, we tested two block-splitting schemes. Given S workers, DSGLD-S ('S' stands for square) partitions $\mathbf{R}$ into $\sqrt{S} \times \sqrt{S}$ blocks as in Fig. 1 (a), i.e. DSGLD-S tries to maximize the within-chain parallelism by using as many orthogonal blocks as possible. We run $\sqrt{S}$ parallel chains, where each chain updates $\sqrt{S}$ sub-parameter blocks in parallel using $\sqrt{S}$ workers. Therefore, all chains can update all parameter at every round. The second splitting scheme, called DSGLD-C ('C' stands for column blocks) divides $\mathbf{R}$ into S blocks as shown in Fig. 1(b). We split $\mathbf{R}$ along the rows because in our experiments we have many more users than items. The blocks in DSGLD-C are not orthogonal because all columns are shared, so we just run S independent parallel chains.

For Gibbs sampling, we use the original BPMF model[3] described in section 2.1. For the other algorithms, we slightly extend the model described in section 3.1 (as in [9, 17]). The extension includes user and item specific bias terms a_i and b_j respectively so that the predictions are modeled as:

$$R_{ij} \approx U_i^\top V_j + a_i + b_j \qquad (31)$$

We use the following priors and hyper-priors for a_i and b_j:

$$a_i \sim \mathcal{N}(0, \lambda_a^{-1}), \quad b_j \sim \mathcal{N}(0, \lambda_b^{-1}),$$
$$\lambda_a, \lambda_b \sim \text{Gamma}(\alpha_0, \beta_0).$$

For $\mathbf{U}$ and $\mathbf{V}$, we use the same priors and hyper-priors as described in Section 3.1. Note that, in the new model, we have to sample $a_i, b_j, \lambda_a, \lambda_b$ in addition to $\mathbf{U}, \mathbf{V}, \Lambda_U, \Lambda_V$. The DSGLD update rules for a_i and b_j are:

$$a_{i,t+1} \leftarrow a_{i,t} + \frac{\epsilon_t}{2}\left\{ \frac{N^{(s)}}{v^{(s)}}\bar{g}(a_{i,t}; \mathcal{M}_t^{(s)}) - \frac{\lambda_a a_{i,t}}{\bar{h}_{i*}} \right\} + \nu_t \quad (32)$$

$$b_{j,t+1} \leftarrow b_{j,t} + \frac{\epsilon_t}{2}\left\{ \frac{N^{(s)}}{v^{(s)}}\bar{g}(b_{j,t}; \mathcal{M}_t^{(s)}) - \frac{\lambda_b b_{j,t}}{\bar{h}_{*j}} \right\} + \nu_t. \quad (33)$$

The main goal of our experiments is to answer the following questions:

- *Accuracy*: How does DSGLD compare to other methods in terms of prediction RMSE?

- *Speed*: How fast can DSGLD achieve the RMSE obtained by 1) optimization algorithms (SGD, DSGLD) 2) Gibbs sampling?

- *Factors which affect the above*: The number of workers, number of chains, block splitting schemes and the latent factor dimension.

4.2 Setup

Dataset	# users	# items	# ratings
Netflix	480K	18K	100M
Yahoo	1.8M	136K	700M

Table 2: Datasets.

We compare all 5 algorithms on two large datasets, Netflix movie ratings [7] and Yahoo music ratings [3] (details in Table 2). To the best of our knowledge, the Yahoo dataset was one of the largest publicly available datasets when we performed the experiments. Note that the Yahoo dataset we use here is different from the one

---

[3]Using the simplified model does not reduce the computation complexity of the Gibbs sampling.

used in the KDD'11 Cup [11] (which has $\sim$250M music ratings and is often referred to by the same name). For the Netflix dataset, we use 80% of the ratings for training and the remaining 20% for testing as in [10]. For the Yahoo dataset, the memory footprint was around 17GB for the train and test ratings, and around 1GB for $\mathbf{U}$ and $\mathbf{V}$ with $D = 60$ in our 64-bit float based implementation. The memory footprint of the Netflix dataset was relatively small.

We used Julia [8] to configure the cluster and execute the core routines of the algorithms. The core routines were implemented in C for high performance. For distributed computing, we used Amazon EC2 instances [1] of type "r3" which were equipped with Intel Xeon 2.5 GHz CPUs and had memory configurable up to 244GB. Although the instances had multiple cores, we restricted all algorithms, except Gibbs sampling, to run on a single-core. For Gibbs sampling, we used a 12-core machine with the same CPU speed. All algorithms were implemented as an in-memory execution model and thus no disk I/O overheads were considered.

We annealed the step size according to the schedule $\epsilon_t = \epsilon_0(1 + t/\kappa)^{-\gamma}$, (as in [5, 23]) which satisfies the convergence conditions in Eqn. (11). We found κ, which controls the decay rate, over the range $\kappa = [10, 50, 100, 500, 1000, 1500]$. The initial step size ϵ_0 was also selected from [9e-6,1e-6] for Netflix and [3e-6,8e-7] for Yahoo. More detailed settings are given in the Appendix. We decreased the stepsize after every *round* which we set to 50 updates. We used $\gamma = 0.51$ in all experiments.

We set the hyperparameters $\tau = 2.0$ and $\alpha_0 = 1.0$ for all experiments. We used $\beta_0 = 1.0$ for all algorithms except SGLD and DSGLD. For SGLD and DSGLD, the scale of the prior gradients sometimes became large due to multiplication by the bias correctors $1/h_{i*}$ and $1/h_{*j}$. In this case, instead of increasing the mini-batch size to reduce the scale of the correctors, we used a more appropriate scale parameter for the Gamma prior distribution ($\beta_0 = 300$), to stabilize the scale of precisions sampled from the posterior Gamma distribution.

Mini-batch sizes were set to 50K data points for Netflix and 100K for Yahoo. The initial values for the precisions Λ were all chosen to be 2.0 after testing over a range $[10, 5, 2, 1, 0.1, 0.01]$. In SGLD and DSGLD, the precision parameters were sampled every 50 rounds after burn-in. We discarded (burned) samples until the RMSE reached 0.85 for Netflix and 1.08 for Yahoo. For DSGLD, which deploys multiple chains, we used the arithmetic mean of the RMSE of all chains to determine whether burn-in has completed. We set the thinning interval to 10 rounds, i.e. we use only every 10^{th} sample to compute the average prediction. The Gibbs sampler in our experiments was initialized near a MAP state which we found using SGD during burn-in.

Running DSGLD requires a block scheduler (line 4 in Algorithm 2) that determines which blocks (workers) are used by each chain in a round. In our experiments, the blocks and the orthogonal groups were chosen beforehand and were assigned to chains deterministically using a cyclic-shift (rotation) at every round with equal visiting frequency. This scheduling policy is illustrated in Fig. 2.

4.3 Results

4.3.1 Convergence and wall-clock time

We first compare the RMSE of the algorithms as a function of computational time. In this experiment, we set $D=30$ for both datasets and used 9 workers for Netflix and 16 workers for Yahoo. Given S workers, we used a $\sqrt{S} \times \sqrt{S}$ block-split for DSGLD-S, $S \times 1$ split for DSGLD-C and $S \times S$ split for DSGD. The total runtime was set to 50K seconds ($\approx$14 hours) for Netflix and 100K seconds ($\approx$27 hours) for Yahoo. In both Figs. 3 and 4, the x-axis
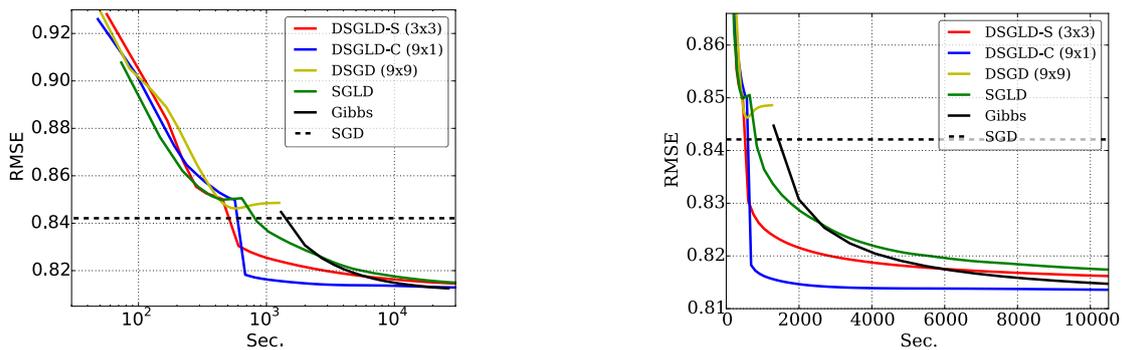
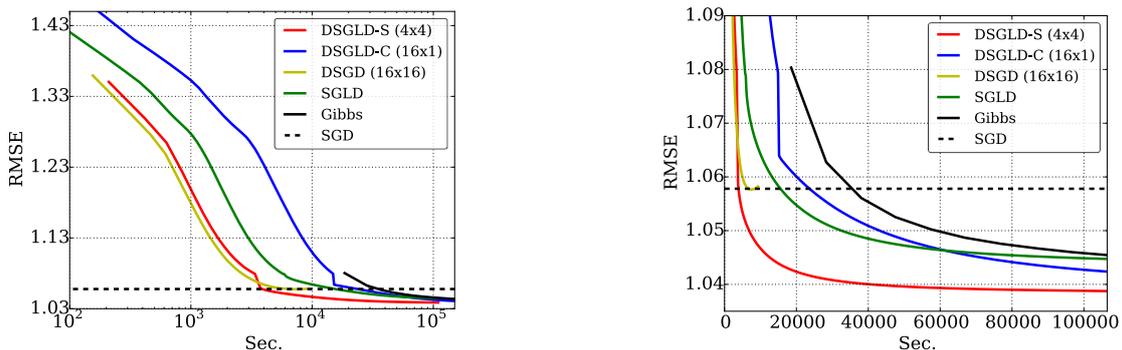**Figure 3: Netflix dataset ($D = 30$).**

**Figure 4: Yahoo Music Rating dataset ($D = 30$).**

is in log-scale for the figure on the left and in linear-scale for the figure on the right.

In Fig. 3, we show results on the Netflix dataset (which is smaller than the Yahoo dataset). We see that in the early (burn-in) stage, all algorithms except Gibbs reduce error at a similar rate. Even though DSGLD-S and DSGD uses block orthogonality to update the sub-parameters of a chain in parallel, because of communication overheads, the gain in speed-up is not enough to outperform a non-distributed algorithm like SGLD which is able to reduce the error at a similar rate (because the dataset size is not very large) without any communication overhead. Note that because there are many chains for DSGLD, we plot the RMSE from only one chain during burn-in. The variance of RMSE across the chains was small during burn-in.

When the burn-in phase ends at around 500 - 700 seconds, MCMC algorithms (SGLD, DSGLD, and Gibbs) begin to collect samples and average their predictions over the samples, while DSGD does not and begins to overfit. Interestingly, at this point, we see a remarkably steep decrease in error for both DSGLD-S and DSGLD-C. In particular, we see the largest decrease for DSGLD-C which deploys 9 independent chains (whereas DSGLD-S uses 3 chains). Note that this is not solely a consequence of collecting a larger number of samples from multiple chains. We believe that the averaged prediction using many independent chains provides better generalization because many modes are likely to be explored (or, a large area of a single broad mode can be covered quickly if many chains reside there). After more investigation, we indeed observed that the same number of samples collected from a single chain (e.g. SGLD) cannot achieve the same level of accuracy obtained with multiple randomly initialized chains. Furthermore, we observed that given a lot more computational time, SGLD and DSGLD-S

can approach the RMSE obtained by DSGLD-C as they also get a chance to explore other modes or to cover a larger area of a single mode. We will revisit the effect of multiple chains in more detail in the next section. Finally, note that Gibbs sampling achieves lower RMSE than DGSLD-C after around 20K seconds (5.5 hours) as shown in Fig. 3 left (but the difference to DSGLD-C is very small). Note that for this dataset, D and $L + M$ were not too large and we used 12-core single machine for parallel Gibbs sampling. Therefore the computational cost of each iteration was not extremely high.

We present our results on the Yahoo dataset in Fig. 4 with $S = 16$ workers. A remarkable point is that, here, unlike with the Netflix dataset, DSGLD-S outperforms DSGLD-C. This is because using orthogonal blocks increases the number of parameters updated per round, resulting in increased convergence speed even after offsetting the communication overhead. As expected, a similar effect is observed for DSGD. The progress of parameter updates in DSGLD-C is relatively slow, requiring $S = 16$ rounds to update all the parameters. Besides, DSGLD-C has a much larger communication overhead because the full matrix $\mathbf{V}$ has to be transferred between the parameter server and each of the workers, whereas only a small block of $\mathbf{V}$ is transferred in DSGLD-S. Specifically, in DSGLD-C the parameter server sends and receives packets of total size $\mathcal{O}((L + SM)D)$ per round whereas in DSGLD-S the total packet size is only $\mathcal{O}((L + M)D)$. Although DSGLD-C is rather slow during burn-in, after burn-in we still see a faster decrease in RMSE compared to SGLD because multiple chains can mix better. Gibbs sampling converges slower than it does on the Netflix dataset because for the Yahoo dataset the number of latent vectors to update, i.e. $L + M$, increases by a factor of four, and the number of ratings, N, by a factor of seven.
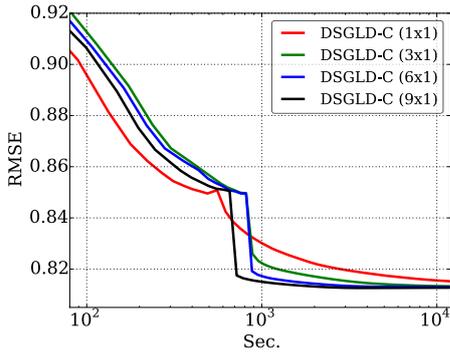
(a) DSGLD-C on the Netflix dataset
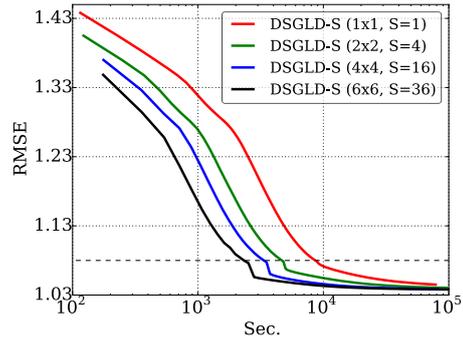
(b) DSGLD-S on the Yahoo dataset

Figure 5: The effect of the number of chains, number of workers, and block split.

For the Netflix dataset, after 1K seconds, DSGLD-C achieved the RMSE (0.8145) that the Gibbs sampler obtains at 10K seconds. Similarly, after 11K seconds, DSGLD-S achieved the RMSE (1.0454) that the Gibbs sampler obtains at 100K seconds. Therefore, the proposed method converges an order of magnitude faster than Gibbs sampling on both datasets, which is especially important when we only have a limited computational budget.

DSGD converges to a prediction RMSE of 0.8462 on Netflix and 1.0576 on Yahoo after 1K seconds and 10K seconds respectively. Given the same amount of computational time, DSGLD achieves an error of 0.8161 on Netflix and 1.0465 on Yahoo, a relative improvement of 3.7% and 1.1%. After convergence, the RI increases to 4.1% for Netflix and 1.8% for Yahoo (See Table. 3).

4.3.2 Number of chains and workers

We also investigated the effect of the number of chains and the number of workers. The results are presented in Fig. 5. According to the observations from the previous experiment, we used DSGLD-C for Netflix and DSGD-S for Yahoo to study this effect. The latent feature dimension was set to D=30.

In Fig. 5 (a), we compare DSGLD-C with $[1, 3, 6, 9]$ chains (and workers) and in each case we evenly split the rows of the rating matrix between the chains. Note that DSGLD-C (1x1) is the same as SGLD running on a single-machine. We see that during burn-in DSGLD-C (1x1) converges faster than the other splits because there is no communication overhead. After burn-in, when the chains start averaging predictions, we see a sharp decrease in error for the other splits. Although splits with more chains decrease error much faster, they all eventually converge to a similar value. Due to poor mixing, a single chain (i.e. SGLD) converges very slowly.

In Fig. 5 (b), we show results for DSGLD-S on the Yahoo dataset. We increased the number of workers to $[1, 4, 16, 36]$ to compare $[1, 2, 4, 6]$ parallel chains. Again DSGLD-S (1x1) denotes SGLD running on a single machine. We see that SGLD converges much more slowly because the dataset is larger than Netflix and SGLD has to update more parameters sequentially. Using more orthogonal blocks, DSGLD-S can update more parameters in parallel and we see more speed-up as we increase the number of workers. Although we increase the number of workers quadratically between the experiments, the packet size transferred between the parameter server and the workers stays constant at $\mathcal{O}((L+M)D)$ because the block size also reduces accordingly. Even after burn-in (horizontal dotted black line at 1.08 RMSE) we see that with more chains we can decrease the error faster. This is because (i) multiple chains help to mix better by exploring a broader space (ii) each chain can

mix faster by updating orthogonal blocks in parallel.

4.3.3 Latent feature dimension

In Fig. 6, we show how the latent feature dimension affects the final RMSE. The final RMSE on Netflix is measured after 50K seconds (14 hours) of computational time, because by then all algorithms had converged (except Gibbs sampling which is expected to take much longer). On the Yahoo dataset, we increased the computational time to 100K secs (1 day), 200K secs (2.3 days), and 300K secs (3.5 days) for D=[30,60,100], respectively, to give the Gibbs sampler more time to converge. In table 3, we show the RMSEs of the different algorithms and the *relative* improvement (or deterioration) compared to DSGLD. The Relative Improvement (RI) of an algorithm x is defined as $RI(x) = (r_d - r_x)/r_d$, where r_x is the RMSE achieved by algorithm x and r_d is the RMSE obtained using DSGLD.

In both Fig. 6 (a) and (b), we see a large difference in performance between SG-MCMC (SGLD and DSGLD) and the optimization methods (SGD and DSGD). The RI is $3.6\% - 4.6\%$ on Netflix and $1.8\% - 3.9\%$ on the Yahoo dataset. As observed in [27], we see that the optimization methods do not consistently improve with increasing D. One reason is that optimization methods are highly sensitive to the hyperparameter values which become difficult to tune as the model becomes more complex. However, our method consistently improves as we increase D, because the hyperparameters are sampled from their posterior distributions. We also see that the performance of Gibbs sampling on Netflix gets worse as D increases, because we used the same amount of computational budget for all D although the computation complexity increases as D does. On the Yahoo dataset on which we increase computational time as we increase D, we see that the RMSE for Gibbs increases as D increases, but is still lower than that of DSGLD.

In Fig. 6 (c), we compare the time (in seconds) required to draw a single sample for the three sampling algorithms at different values of D on the Yahoo dataset. We see that Gibbs sampling is almost two orders of magnitude slower than SGLD. For D=100, SGLD, DSGLD-S, and Gibbs generated 688, 460, and 8 samples respectively in 300K seconds of computational time. For Netflix, Gibbs generated around 100 samples in 50K seconds for D=30. Thus, even though the Gibbs sampler can produce higher quality samples (in terms of lower auto-correlation), the sampling speed is so slow that it cannot satisfactorily handle large scale datasets.

5. CONCLUSION

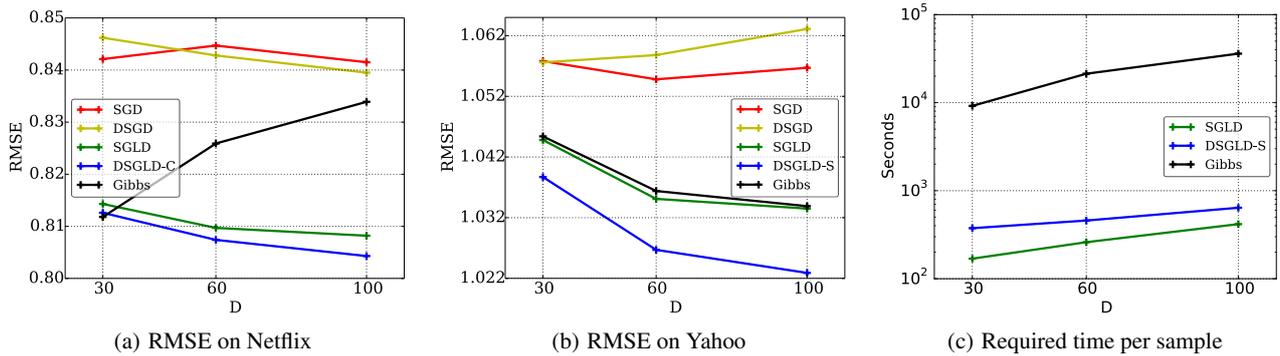| (a) RMSE on Netflix | (b) RMSE on Yahoo | (c) Required time per sample |

Figure 6: The effect of the latent feature dimension. (a) and (b) show RMSE for $D = [30, 60, 100]$ on (a) the Neflix dataset and (b) the Yahoo music ratings dataset. The maximum computational time was set to 50K seconds for Netflix and 100K (D=30), 200K (D=60), and 300K (D=100) seconds for Yahoo. (c) shows time (in seconds) required to draw a single sample on the Yahoo dataset.

D	SGD	DSGD	SGLD	DSGLD-C	Gibbs		D	SGD	DSGD	SGLD	DSGLD-S	Gibbs
30	0.8421	0.8462	0.8143	0.8126	**0.8118**		30	1.0578	1.0576	1.0448	**1.0387**	1.0454
	-3.63%	-4.13%	-0.21%	-	+0.09%			-1.83%	-1.82%	-0.58 %	-	-0.64%
60	0.8447	0.8428	0.8097	**0.8074**	0.8259		60	1.0548	1.0588	1.0351	**1.0267**	1.0364
	-4.62%	-4.38%	-0.28%	-	-2.29%			-2.73%	-3.13%	-0.82%	-	-0.94%
100	0.8415	0.8395	0.8082	**0.8043**	0.8339		100	1.0567	1.0631	1.0335	**1.0229**	1.0339
	-4.63%	-4.37%	-0.48%	-	-3.68%			-3.30%	-3.93%	-1.04%	-	-1.08%

Table 3: RMSE and relative improvement (RI). Left: Netflix. Right: Yahoo. The percentage shown below each RMSE value is the relative improvement.

Most applications of matrix factorization to recommender systems are based on stochastic gradient optimization algorithms because these are the only ones that can computationally handle very large datasets. However, by restricting ourselves to such simple algorithms, we miss out on all the advantages of Bayesian modelling such as quantifying uncertainty, controlling over-fitting, incorporating prior information and better prediction accuracy. In this paper, we introduced a novel algorithm for scalable distributed Bayesian matrix factorization that achieves the best of both worlds, i.e. it inherits all the advantages of Bayesian inference at the speed of stochastic gradient optimization.

Our algorithm, based on Distributed Stochastic Gradient Langevin Dynamics, uses only a mini-batch of ratings to make each update as in Stochastic Gradient Descent optimization. By running multiple chains in parallel, and also using multiple workers within a chain to update orthogonal blocks, we can scale up Bayesian Matrix Factorization to very large datasets. Parallel chains with different random initializations also help us to average predictions from multiple modes and improve accuracy. Moreover, our algorithm can effectively handle datasets that are distributed across multiple machines unlike traditional MCMC algorithms.

We believe that our method is just one example of a much larger class of scalable distributed Bayesian matrix factorization methods. For example, we can consider using more sophisticated stochastic gradient algorithms [5, 23, 9, 10] in place of SGLD to further improve the mixing rate.

Acknowledgments

References

[1] Amazon ec2 instances. http://aws.amazon.com/ec2/instance-types/.

[2] Netflix prize. http://www.netflixprize.com.

[3] R2 - yahoo! music user ratings of songs with artist, album, and genre meta information, v. 1.0 (1.4 gbyte and 1.1 gbyte). http://webscope.sandbox.yahoo.com/.

[4] R. Adams, G. Dahl, and I. Murray. Incorporating side information in probabilistic matrix factorization with gaussian processes. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*, 2010.

[5] S. Ahn, A. Korattikara, and M. Welling. Bayesian posterior sampling via stochastic gradient fisher scoring. In *International Conference on Machine Learning*, 2012.

[6] S. Ahn, B. Shahbaba, and M. Welling. Distributed stochastic gradient mcmc. In *International Conference on Machine Learning (ICML)*, 2014.

[7] J. Bennett and S. Lanning. The netflix prize. *In KDD Cup and Workshop in conjunction with KDD*, 2007.

[8] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah. Julia: A fresh approach to numerical computing. *CoRR*, 2014. http://dblp.uni-trier.de/rec/bib/journals/corr/BezansonEKS14.

[9] T. Chen, E. Fox, and C. Guestrin. Stochastic gradient hamiltonian monte carlo. In *International Conference on Machine Learning (ICML)*, 2014.

[10] N. Ding, Y. Fang, R. Babbush, C. Chen, R. Skeel, and H. Neven. Bayesian sampling using stochastic gradient thermostats. In *Advances in Neural Information Processing Systems (NIPS)*, 2014.

[11] G. Dror, N. Koenigstein, Y. Koren, and M. Weimer. The yahoo! music dataset and kdd-cup'11. In *Proceedings of KDD-Cup 2011 competition*, 2012.

[12] S. Duane, A. Kennedy, B. Pendleton, and D. Roweth. Hybrid monte carlo. *Physics letters B*, 195(2):216–222, 1987.

[13] R. Gemulla, E. Nijkamp, P. Haas, and Y. Sismanis. Large-scale matrix factorization with distributed stochastic gradient descent. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2011.

[14] M. Girolami and B. Calderhead. Riemann manifold langevin and hamiltonian monte carlo. *Journal of the Royal Statistical Society B*, 73 (2):1–37, 2010.

[15] K. B. Hall, S. Gilpin, and G. Mann. Mapreduce/bigtable for distributed optimization. In *NIPS LCCC Workshop*, 2010.

[16] A. Korattikara, Y. Chen, and M. Welling. Austerity in mcmc land: Cutting the metropolis-hastings budget. In *International Conference on Machine Learning (ICML)*, 2014.

[17] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. In *IEEE Computer*, 2009.

[18] G. Mann, R. McDonald, M. Mohri, N. Silberman, and D. Walker. Efficient large-scale distributed training of conditional maximum entropy models. In *Neural Information Processing Systems*, 2009.

[19] R. McDonald, K. Hall, and G. Mann. Distributed training strategies for the structured perceptron. In *HLT*, 2010.

[20] A. Mnih and R. Salakhutdinov. Probabilistic matrix factorization. In *Advances in Neural Information Processing Systems*, 2007.

[21] R. Neal. Mcmc using hamiltonian dynamics. In S. Brooks, A. Gelman, G. Jones, and X. Meng, editors, *Handbook of Markov Chain Monte Carlo*. Chapman&Hall/CRC, 2011.

[22] F. Niu, B. Recht, C. Ré, and S. J. Wright. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *arXiv preprint arXiv:1106.5730*, 2011.

[23] S. Patterson and Y. W. Teh. Stochastic gradient riemannian langevin dynamics on the probability simplex. In *Advances in Neural Information Processing Systems*, 2013.

[24] I. Porteous, A. Ascuncion, and M. Welling. Bayesian matrix factorization with side information and dirichlet process mixtures. In *AAAI Conference on Artificial Intelligence*, 2010.

[25] B. Recht and C. Re. Parallel stochastic gradient algorithms for large-scale matrix completion. In *Mathematical Programming Computation*, 2013.

[26] P. Rossky, J. Doll, and H. Friedman. Brownian dynamics as smart monte carlo simulation. In *The Journal of Chemical Physics*, 1978.

[27] R. Salakhutdinov and A. Mnih. Bayesian probabilistic matrix factorization using markov chain monte carlo. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, 2008.

[28] C. Teflioudi, F. Makari, and R. Gemulla. Distributed matrix completion. In *IEEE 12th International Conference on Data Mining*, 2012.

[29] M. Welling and Y. W. Teh. Bayesian learning via stochastic gradient langevin dynamics. In *International Conference on Machine Learning (ICML)*, 2011.

[30] Y. Zhuang, W. S. Chin, Y. C. Juan, and C. J. Lin. A fast parallel sgd for matrix factorizatio in shared memory systems. In *Proceedings of the 7th ACM conference on Recommender systems*, 2013.

[31] M. Zinkevich, M. Weimer, and A. Smola. Parallelized stochastic gradient descent. In *Neural Information Processing Systems*, 2010.

APPENDIX

A.   STEP-SIZE PARAMETERS

	SGD	DSGD	SGLD	DSGLD-C	DSGLD-S
ϵ_0	9e-6	1e-6	9e-6	9e-6	3e-6
κ	50	10	1000	1000	500

Table 4: Stepsize parameters for Netflix D=30 and 9 workers

	SGD	DSGD	SGLD	DSGLD-C	DSGLD-S
ϵ_0	1.5e-6	3e-7	1.5e-6	9e-7	1.5e-6
κ	500	100	1000	1000	500

Table 5: Stepsize parameters for Yahoo D=30 and 16 workers