



# Simulation comparisons between Bayesian and de-biased estimators in low-rank matrix completion

The Tien Mai<sup>1</sup>

Received: 8 August 2022 / Accepted: 7 January 2023  
© The Author(s) 2023

## Abstract

In this paper, we study the low-rank matrix completion problem, a class of machine learning problems, that aims at the prediction of missing entries in a partially observed matrix. Such problems appear in several challenging applications such as collaborative filtering, image processing, and genotype imputation. We compare the Bayesian approaches and a recently introduced de-biased estimator which provides a useful way to build confidence intervals of interest. From a theoretical viewpoint, the de-biased estimator comes with a sharp minimax-optimal rate of estimation error whereas the Bayesian approach reaches this rate with an additional logarithmic factor. Our simulation studies show originally interesting results that the de-biased estimator is just as good as the Bayesian estimators. Moreover, Bayesian approaches are much more stable and can outperform the de-biased estimator in the case of small samples. In addition, we also find that the empirical coverage rate of the confidence intervals obtained by the de-biased estimator for an entry is absolutely lower than of the considered credible interval. These results suggest further theoretical studies on the estimation error and the concentration of Bayesian methods as they are quite limited up to present.

**Keywords** Matrix completion · Bayesian method · De-biased estimator · Uncertainty quantification · Confidence interval

## 1 Introduction

The goal of low-rank matrix completion is to recover a low-rank matrix from its partially (noisy) observed entries. This problem has recently received an increased attention due to the emergence of several challenging applications, such as recommender systems [1, 38] (particularly the famous Netflix challenge [6]), genotype imputation [12, 19], image processing [7, 18, 25] and quantum state tomography [15, 28, 30]. Different approaches from frequentist

---

✉ The Tien Mai  
the.t.mai@ntnu.no

<sup>1</sup> Department of Mathematical Sciences, Norwegian University of Science and Technology, Trondheim, Norway

to Bayesian methods have been proposed and studied from theoretical and computational points of views, see for example [2, 3, 8–11, 21, 23, 24, 26, 29, 34, 37, 40].

From a frequentist point of view, most of the recent methods are usually based on penalized optimization. A seminal result can be found in Candés and Recht [9], Candés and Tao [10] for exact matrix completion (noiseless case) and further developed in the noisy case in Candés and Plan [8], Koltchinskii et al. [21], Negahban and Wainwright [33]. Some efficient algorithms had also been studied, for example see [17, 31, 34]. More particularly, in the notable work [21], the authors studied nuclear-norm penalized estimators and provided reconstruction error rate for their methods. They also showed that these error rates are minimax-optimal (up to a logarithmic factor). Note that the error rate, i.e. the average quadratic error on the entries, of a rank- $r$  matrix size  $m \times p$  from  $n$ -observations can not be better than:  $r \max(m, p)/n$  [21].

More recently, in a work by Chen et al. [11], de-biased estimators have been proposed for the problem of noisy low-rank matrix completion. The estimation accuracy of this estimator is shown to be sharp in the sense that it reaches the minimax-optimal rate without any additional logarithmic factor. A sharp bound has also been obtained by a different estimator in Klopp [20]. However, uncertainty quantification is not given. More importantly, the confidence intervals on the reconstruction of entries of the underlying matrix are also provided by using the de-biased estimators in the work by Chen et al. [11]. It is noted that conducting uncertainty quantification for matrix completion is not straightforward. This is because, in general, the solutions for matrix completion do not admit closed-form and the distributions of the estimates returned by the state-of-the-art algorithms are hard to derive.

On the other hand, uncertainty quantification can be obtained straightforwardly from a Bayesian perspective. More specifically, the unknown matrix is considered as a random variable with a specific prior distribution and statistical inference can be obtained using the posterior distribution, for example considering credible intervals. Bayesian methods have been studied for low-rank matrix completion mainly from a computational viewpoint [4, 4, 5, 13, 23, 24, 37, 39, 40]. Most Bayesian estimators are based on conjugate priors which allow to use Gibbs sampling [4, 37] or Variational Bayes methods [24]. These algorithms are fast enough to deal with and actually tested on large datasets like Netflix [6] or MovieLens [16]. However, the theoretical understanding of Bayesian estimators is quite limited, up to our knowledge, [29] and [3] are the only prominent examples. More specifically, they showed that a Bayesian estimator with a low-rank factorization prior reaches the minimax-optimal rate up to a logarithmic factor and the paper [3] further shows that the same estimation error rate can be obtained by using a Variational Bayesian estimator.

In this paper, to understand the performance of Bayesian approaches when compared to the de-biased estimators, we perform numerical comparisons on the estimation accuracy (the estimation error, the normalized squared error and the prediction error, see Sect. 3) considering the de-biased estimator in Chen et al. [11] and the Bayesian methods [3] for which the statistical properties have been well studied. Furthermore, we examine in detail the behaviour of the confidence intervals obtained by the de-biased estimator and the Bayesian credible intervals. Interestingly, it is noted that recent works [35, 36] show that Bayesian methods are now the most accurate in practical recommender systems. Although Bayesian methods have become popular in the problem of matrix completion, its uncertainty quantification (e.g. credible intervals) has received much more limited attention in the literature.

Results from simulation comparisons release originally interesting messages. More specifically, the de-biased estimator is just as good as the Bayesian estimators when we look at the estimation accuracy, although it is completely successful in improving the estimator being de-biased. On the other hand, the Bayesian approaches are much more stable than the de-biased method and, in addition, they outperform the de-biased estimator especially in the

case of small samples. Moreover, we find that the coverage rates of the 95% confidence intervals obtained using the de-biased estimator are lower than the 89% equal-tailed credible intervals. These evidences suggest that the Bayesian estimators may actually reach the minimax-optimal rate sharply and the log-term could be due to the technical proofs (the PAC-Bayesian bounds technique). Furthermore, the concentration rate of the corresponding Bayesian posterior discussed in Alquier and Ridgway [3] with a log-term might not be tight.

The rest of the paper is structured as follows. In Sect. 2 we present the low-rank matrix completion problem, then introduce the de-biased estimator and the corresponding confidence interval and provide details on the considered Bayesian estimators. In Sect. 3, simulation studies comparing the different methods are presented. We discuss our results and give some concluding remarks in the final section.

## 2 Low-rank matrix completion

### 2.1 Model

In this work, we adopt the statistical model commonly studied in the literature for noisy matrix completion [11]. Let  $M^* \in \mathbb{R}^{m \times p}$  be an unknown rank- $r$  matrix of interest. We partially observe some noisy entries of  $M^*$  as

$$Y_{ij} = M_{ij}^* + \mathcal{E}_{ij}, \quad (i, j) \in \Omega \tag{1}$$

where  $\Omega \subseteq \{1, \dots, m\} \times \{1, \dots, p\}$  is a small subset of indexes and  $\mathcal{E}_{ij} \sim \mathcal{N}(0, \sigma^2)$  are independently generated noise at the location  $(i, j)$ . The random sampling model is assumed that each index  $(i, j) \in \Omega$  is observed independently with probability  $\kappa$  (i.e., data are missing uniformly at random). Then, the problem of estimating  $M^*$  with  $n = |\Omega| < mp$  is called the (noisy) low-rank matrix completion problem.

Let  $\mathcal{P}_\Omega(\cdot) : \mathbb{R}^{m \times p} \mapsto \mathbb{R}^{m \times p}$  be the orthogonal projection onto the observed entries in the index set  $\Omega$  that

$$\mathcal{P}_\Omega(Y)_{ij} = \begin{cases} Y_{ij}, & \text{if } (i, j) \in \Omega, \\ 0, & \text{if } (i, j) \notin \Omega \end{cases} .$$

**Notations:** For a matrix  $A \in \mathbb{R}^{m \times p}$ ,  $\|A\|_F = \sqrt{\text{trace}(A^\top A)}$  denotes its Frobenius norm and  $\|A\|_* = \text{trace}(\sqrt{A^\top A})$  denotes its nuclear norm.  $[a \pm b]$  denotes the interval  $[a - b, a + b]$ . We use  $I_q$  to denote the identity matrix of dimension  $q \times q$ .

### 2.2 The de-biased estimator

Let  $\hat{M}$  be either the solution of the following nuclear norm regularization [31]

$$\min_{Z \in \mathbb{R}^{m \times p}} \frac{1}{2} \|\mathcal{P}_\Omega(Z - Y)\|_F^2 + \lambda \|Z\|_*,$$

or of the following factorization minimization [17]

$$\min_{U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{p \times r}} \frac{1}{2} \|\mathcal{P}_\Omega(Y - UV^\top)\|_F^2 + \frac{\lambda}{2} \|U\|_F^2 + \frac{\lambda}{2} \|V\|_F^2, \tag{2}$$

where  $\lambda > 0$  is a tuning parameter. The optimization problem in (2) can be seen as the problem of finding the MAP (maximum a posteriori) in Bayesian modeling where Gaussian priors are used on columns of the factors  $U$  and  $V$ , detailed discussion can be found in d in Alquier et al. [4], Fithian and Mazumder [14].

Given an estimator  $\hat{M}$  as above, the de-biased estimator [11] is defined as

$$M^{db} := \Pr_{\text{rank-}r} \left[ \hat{M} - \mathcal{P}_{\Omega}(\hat{M} - Y) \right], \quad (3)$$

where  $\Pr_{\text{rank-}r}(B) = \arg \min_{A, \text{rank}(A) \leq r} \|A - B\|_F$  is the projection onto the set of rank- $r$  matrices.

**Remark 1** The estimation accuracy of the de-biased estimator, provided in Theorem 3 in Chen et al. [11] under some assumptions, is  $\|M^{db} - M^*\|_F^2 \leq c \max(m, p)r\sigma^2/n$  without any extra log-term and  $c$  is universal numerical constant.

### 2.2.1 Confidence interval

Let  $\hat{M} = \hat{U} \hat{\Sigma} \hat{V}^\top$  be the singular values decomposition of  $\hat{M}$ . Put

$$v_{ij} := \sigma^2 \left[ U_{i \cdot}^{db} (U^{db\top} U^{db})^{-1} U_{i \cdot}^{db\top} + V_{j \cdot}^{db} (V^{db\top} V^{db})^{-1} V_{j \cdot}^{db\top} \right] / \kappa, \quad (4)$$

where

$$U^{db} = \hat{U}(\hat{\Sigma} + (\lambda/\kappa)I_r)^{1/2} \text{ and } V^{db} = \hat{V}(\hat{\Sigma} + (\lambda/\kappa)I_r)^{1/2}.$$

Then, given a significance level  $\alpha \in (0, 1)$ , the following interval

$$\left[ M_{ij}^{db} \pm \Phi^{-1}(1 - \alpha/2) \sqrt{v_{ij}} \right]$$

is a nearly accurate two-sided  $(1 - \alpha)$  confidence interval of  $M_{ij}^*$ , where  $\Phi(\cdot)$  is the CDF of the standard normal distribution. This is given in Corollary 1 in Chen et al. [11]. This method is implemented in the R package `dbMC` [27].

### 2.3 The Bayesian estimators

The Bayesian estimator studied in Alquier and Ridgway [3] is given by

$$M^B := \int M \rho_\lambda(M|Y) dM$$

where

$$\rho_\lambda(M|Y) \propto L(Y|M)^\lambda \pi(M)$$

is the posterior and  $L(Y|M)^\lambda$  is the likelihood raised to the power  $\lambda$ . Here  $\lambda \in (0, 1)$  is a tuning parameter and  $\pi(M)$  is the prior distribution.

### 2.3.1 Priors

A popular choice for the priors in Bayesian matrix completion is to assign conditional Gaussian priors to  $U \in \mathbb{R}^{m \times K}$  and  $V \in \mathbb{R}^{p \times K}$  such that

$$M = UV^\top = \sum_{k=1}^K U_{\cdot k} V_{\cdot k}^\top,$$

for a fixed integer  $K \leq \min(m, p)$ . More specifically, for  $k \in \{1, \dots, K\}$ , independently

$$\begin{aligned} U_{\cdot k} &\sim \mathcal{N}(0, \gamma_k I_m), \\ V_{\cdot k} &\sim \mathcal{N}(0, \gamma_k I_p), \\ \gamma_k^{-1} &\sim \Gamma(a, b), \end{aligned} \quad (5)$$

where  $I_q$  is the identity matrix of dimension  $q \times q$  and  $a, b$  are some tuning parameters. This type of prior is conjugate so that the conditional posteriors can be derived explicitly in closed form and allows to use the Gibbs sampler, see [37] for details. Some reviews and discussion on low-rank factorization priors can be found in Alquier [2], Alquier et al. [4].

**Remark 2** In the case that the rank  $r$  is not known, it is natural to take  $K$  as large as possible, e.g.  $K = \min(m, p)$  but this may be computationally prohibitive if  $K$  is large.

**Remark 3** The estimation error for this Bayesian estimator, under some assumptions, given in Corollary 4.2 in Alquier and Ridgway [3], is  $\|M^B - M^*\|_F^2 \leq \max(m, p)r\sigma^2/n$  with an additional (multiplicative) log-term by  $\log(n \max(m, p))$ . It is noted that the rate is also reached in Mai and Alquier [29] with an additional (multiplicative) log-term by  $\log(\min(m, p))$  under general sampling distribution however the authors considered some truncated priors.

For a given rank- $r$ , we propose to consider the following prior, called fixed-rank-prior,

$$\begin{aligned} U_{\cdot k} &\sim \mathcal{N}(0, I_m), \\ V_{\cdot k} &\sim \mathcal{N}(0, I_p), \end{aligned} \quad (6)$$

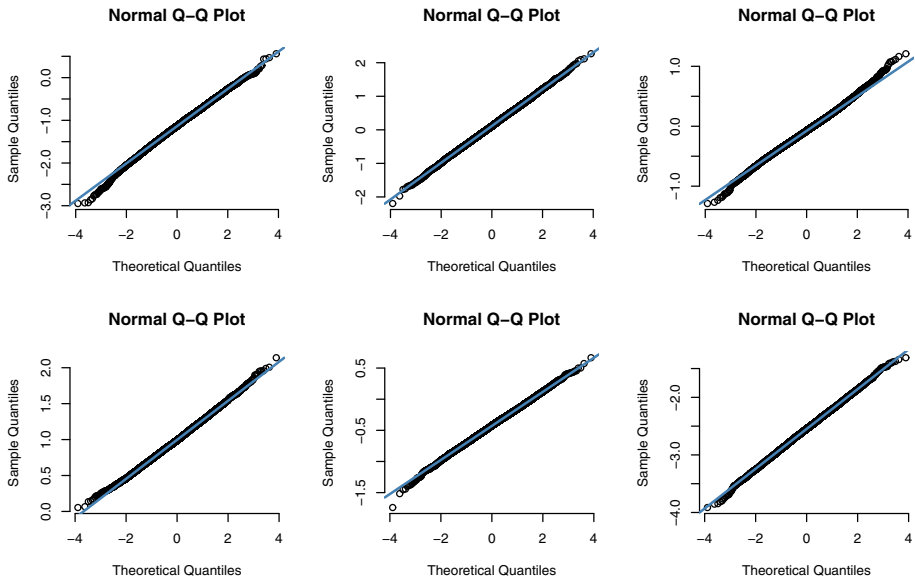
for  $k = 1, \dots, r$ . This prior is a simplified version of the above prior. We note that for  $K > r$  the Gibbs sampler of the fixed-rank-prior will be faster than Gibbs sampler for the above prior. Interestingly, results from simulation for the Bayesian estimator with this prior are slightly better than the one based on the above prior at some point.

**Remark 4** We remark that the theoretical estimation error for the Bayesian estimator with the fixed-rank-prior given in (6) remains unchanged following by Corollary 4.2 in Alquier and Ridgway [3].

### 2.3.2 Credible intervals

Using Bayesian approach, the credibility intervals for the matrix and their functions (e.g. entries) can be easily constructed using the Markov Chain Monte Carlo (MCMC) technique. Here, we focus on the equal-tailed credible interval for an entry (Fig. 1).

More precisely, the credible intervals are reported using the 89% equal-tailed intervals that are recommended by Kruschke [22], McElreath [32] for small posterior samples as in our situations with 500 posterior samples. We noted that, according to Salakhutdinov and Mnih [37] as the data are too big to draw a reasonable size sample, the authors state that drawing



**Fig. 1** Q–Q (quantile–quantile) plot to compare the 10,000 posterior samples for some entries against the standard normal distribution. Top row (from left to right, 3 figures) is the results from Setting I with  $r = 2$ ,  $p = 100$ ,  $\tau = 50\%$ . Bottom row (from left to right, 3 figures) is the results from Setting I with  $r = 2$ ,  $p = 1000$ ,  $\tau = 50\%$

only 500 observations from the Gibbs Sampler took 90 hours for the Netflix dataset. Thus, we focus on the 89% equal-tailed credible intervals for 500 posterior samples. It is, however, noted that to obtain 95% intervals, an effective posterior sample size of at least 10,000 is recommended [22], which is computationally costly to run on all of our simulations. A few examples with 10,000 posterior samples are examined in Fig. 2.

### 3 Simulation studies

#### 3.1 Experimental designs


In order to access the performance of different estimators, a series of experiments were conducted with simulated data. We fix  $m = 100$  and alternate the other dimension by taking  $p = 100$  and  $p = 1000$ . The rank  $r$  is varied between  $r = 2$  and  $r = 5$ .

- Setting I: In the first setting, a rank- $r$  matrix  $M_{m \times p}^*$  is generated as the product of two rank- $r$  matrices,

$$M^* = U_{m \times r}^* V_{p \times r}^{*\top},$$

where the entries of  $U^*$  and  $V^*$  are i.i.d  $\mathcal{N}(0, 1)$ . With a missing rate  $\tau = 20\%$ ,  $50\%$  and  $80\%$ , the entries of the observed set are drawn uniformly at random. This sampled set is then corrupted by noise as in (1), where the  $\mathcal{E}_i$  are i.i.d  $\mathcal{N}(0, 1)$ .

- Setting II: The second series of simulations is similar to the first one, except that the matrix  $M^*$  is no longer rank- $r$ , but it can be well approximated by a rank- $r$  matrix:



**Fig. 2** Plot to compare the limiting Gaussian distributions of the de-biased estimator and the histograms of the 10,000 posterior samples for some entries. The dotted line is the true value of the entries. Top row (from left to right, 3 figures) is the results from Setting I with  $r = 2, p = 100, \tau = 50\%$ . Bottom row (from left to right, 3 figures) is the results from Setting I with  $r = 2, p = 1000, \tau = 50\%$

$$M^* = U_{m \times r}^* V_{p \times r}^{*\top} + \frac{1}{10} A_{m \times 50} B_{p \times 50}^\top$$

where the entries of  $A$  and  $B$  are i.i.d  $\mathcal{N}(0, 1)$ .

- **Setting III:** This setting is similar to Setting I but here a heavy tail noise is used. More specifically, the noise  $\mathcal{E}_i$  are i.i.d Student distribution  $t_3$  with 3 degrees of freedom.
- **Setting IV:** The set up of this setting is also similar to Setting I. However, we consider a more extreme case where the entries of  $U^*, V^*$  and the noise  $\mathcal{E}_i$  are all i.i.d drawn from the Student distribution with 3 degrees of freedom.

**Remark 5** We note that for the second series of simulations, with approximate low-rank matrices, the theory of the de-biased estimator can not be used whereas theoretical guarantees for Bayesian estimators are still valid, see [3]. The setting I follows exactly the minimax-optimal regime and thus it will allow to access the accuracy of the considered estimators. The last 2 settings (III and IV) are misspecification models set up where the theoretical guarantee is not available for all considered estimators.

The behavior of an estimator (say  $\widehat{M}$ ) is evaluated through the average squared error (ase) per entry

$$ase := \frac{1}{mp} \|\widehat{M} - M^*\|_F^2$$

and the relative squared error (rse)

$$rse := \frac{\|\widehat{M} - M^*\|_F^2}{\|M^*\|_F^2}.$$

We also measure the error in predicting the missing entries by using

$$\text{Pred} := \frac{\|\mathcal{P}_{\bar{\Omega}}(\hat{M} - M^*)\|_F^2}{mp - n},$$

where  $\bar{\Omega}$  is the set of un-observed entries. For each setup, we generate 100 data sets (simulation replicates) and report the average and the standard deviation for a measure of error of each estimator over the replicates.

We compare the de-biased estimator (denoted by ‘**d.b**’), the Bayesian estimator with the fixed-rank-prior (6) (denoted by ‘**f.Bayes**’) and the Bayesian estimator with the (flexible rank) prior (5) (denoted by ‘**Bayes**’). As a by-product in calculating the de-biased estimator through the Alternating Least Squares estimator (2), we also report the results for this estimator, denoted it by ‘**als**’.

The ‘als’ estimator is available from the R package ‘softImpute’ [31] and is used with default options. The ‘d.b’ estimator is run with  $\lambda = 2.5\sigma\sqrt{mp}$  as in Chen et al. [11]. The ‘f.Bayes’ and ‘Bayes’ estimators are used with tuning parameter  $\lambda = 1/(4\sigma^2)$  and parameters for the prior of ‘Bayes’ estimator are  $K = 10$ ,  $a = 1$ ,  $b = 1/100$ . The Gibbs samplers for these two Bayesian estimators are run with 500 steps and 100 burn-in steps.

### 3.2 Results on estimation accuracy

From the results in Tables 1 and 2, it is clear that the de-biased estimator significantly outperforms its ancestry estimator being de-biased. Whereas, the de-biased estimator is just as good as the Bayesian methods in some cases.

More specifically, in Table 1, the de-biased estimator behaves similar compared to Bayesian estimators in the case with high rates of observation (say  $\tau = 20\%$  or  $50\%$ ). With the case of highly missing rate  $\tau = 80\%$ , the de-biased estimator returns highly unstable results, this may be because its ancestry estimator (here it is the als estimator) is unstable with few observations. However, when the dimension of the matrix increases, the differences between the de-biased estimator and the Bayesian estimators become smaller. This is also recognized for the setting of approximate low-rank matrices as in Table 2 and in Table 5, 6.

The ‘f.Bayes’ method yields the best results quite often in terms of all considered errors (ase, Nase and Pred) in setting of exact low-rank matrices. However, it is noted that for the setting with the true underlying matrix being approximately low-rank, in Tables 2 and 6, the ‘Bayes’ approach is slightly better than the ‘f.Bayes’ approach at some point. This can be explained as the ‘Bayes’ approach employs a kind of approximate low-rank prior through the Gamma prior on the variance of the factor matrices and thus it is able to adapt to the approximate low-rankness.

Results in the cases of model misspecification with heavy tail noise are given Tables 3 and 4. Although Bayesian methods, especially ‘f.Bayes’ method, yield better results compared with ‘als’ or ‘db’, all methods fail in the case of highly missing data,  $\tau = 80\%$ . This could be due to the fact that these considered methods are all designed for the case of Gaussian noise and thus they are not robust to other heavy tail noise, such as Student noise.





Table 1 continued

		$r = 2, p = 1000, \tau = 20\%$			$r = 5, p = 1000, \tau = 20\%$		
	als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes
Pred	0.029 (.001)	0.029 (.001)	0.029 (.001)	0.029 (.001)	0.076 (.002)	0.076 (.002)	0.076 (.002)
		$r = 2, p = 1000, \tau = 50\%$			$r = 5, p = 1000, \tau = 50\%$		
	als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes
ase	0.523 (.004)	0.046 (.001)	0.046 (.001)	0.046 (.001)	0.564 (.004)	<b>0.119</b> (.002)	0.120 (.002)
rse	0.263 (.026)	0.023 (.002)	0.023 (.002)	0.023 (.003)	0.112 (.007)	0.024 (.002)	0.024 (.001)
Pred	0.048 (.002)	0.048 (.002)	0.048 (.002)	0.048 (.002)	0.129 (.003)	0.129 (.003)	0.129 (.003)
		$r = 2, p = 1000, \tau = 80\%$			$r = 5, p = 1000, \tau = 80\%$		
	als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes
ase	0.344 (.324)	0.167 (.331)	0.135 (.005)	0.136 (.005)	0.594 (.185)	0.450 (.190)	<b>0.408</b> (.012)
rse	0.174 (.172)	0.085 (.175)	0.068 (.007)	0.069 (.007)	0.120 (.040)	0.091 (.041)	0.082 (.011)
Pred	0.180 (.404)	0.180 (.406)	0.141 (.006)	0.142 (.006)	0.492 (.231)	0.474 (.196)	<b>0.440</b> (.014)

The mean and the standard deviation (in parentheses) of each error between the simulation replicates are presented



Table 2 continued

Approximate rank-2, $p = 1000, \tau = 20\%$		Approximate rank-5, $p = 1000, \tau = 20\%$					
	als	d.b	f.Bayes	als	d.b	f.Bayes	Bayes
Pred	0.545 (.012)	0.545 (.012)	0.544 (.012)	0.610 (.014)	0.610 (.014)	0.608 (.014)	0.591 (.011)
Approximate rank-2, $p = 1000, \tau = 50\%$		Approximate rank-5, $p = 1000, \tau = 50\%$					
	als	d.b	f.Bayes	als	d.b	f.Bayes	Bayes
ase	0.786 (.007)	0.546 (.012)	0.545 (.012)	0.846 (.007)	0.624 (.011)	0.621 (.011)	0.619 (.011)
rse	0.315 (.025)	0.219 (.018)	0.219 (.018)	0.155 (.011)	0.115 (.008)	0.114 (.008)	0.114 (.008)
Pred	0.570 (.012)	0.571 (.012)	0.568 (.012)	0.694 (.013)	0.694 (.013)	0.687 (.013)	0.687 (.013)
Approximate rank-2, $p = 1000, \tau = 80\%$		Approximate rank-5, $p = 1000, \tau = 80\%$					
	als	d.b	f.Bayes	als	d.b	f.Bayes	Bayes
ase	0.770 (.013)	0.681 (.015)	0.666 (.014)	1.149 (.079)	1.077 (.081)	0.992 (.021)	0.992 (.021)
rse	0.316 (.027)	0.280 (.024)	0.274 (.023)	0.211 (.018)	0.198 (.018)	0.183 (.011)	0.183 (.011)
Pred	0.712 (.016)	0.712 (.016)	0.693 (.015)	1.186 (.099)	1.185 (.099)	1.080 (.024)	1.082 (.024)

The mean and the standard deviation (in parentheses) of each error between the simulation replicates are presented



Table 3 continued

		$r = 2, p = 100, \tau = 80\%$			$r = 5, p = 100, \tau = 80\%$					
		als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
Pred		3.726 (2.31)	3.732 (2.31)	1.171 (.596)	0.977 (.403)	12.08 (2.38)	12.09 (2.39)	2.830 (.565)	2.682 (.516)	
		$r = 2, p = 1000, \tau = 20\%$			$r = 5, p = 1000, \tau = 20\%$					
		als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
ase		2.464 (.728)	0.171 (.751)	0.120 (.283)	0.283 (.182)	2.558 (.990)	0.368 (1.04)	0.290 (.604)	0.389 (.188)	
rse		1.255 (.415)	0.089 (.404)	0.062 (.151)	0.144 (.096)	0.511 (.194)	0.073 (.202)	0.059 (.129)	0.078 (.039)	
Pred		0.414 (3.10)	0.415 (3.12)	0.177 (.786)	0.293 (.210)	0.825 (4.38)	0.828 (4.40)	0.482 (2.10)	0.435 (.189)	
		$r = 2, p = 1000, \tau = 50\%$			$r = 5, p = 1000, \tau = 50\%$					
		als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
ase		1.633 (.375)	0.236 (.394)	0.213 (.334)	0.399 (.151)	1.978 (1.13)	0.685 (1.16)	0.526 (.630)	0.623 (.240)	
rse		0.824 (.217)	0.120 (.208)	0.109 (.180)	0.202 (.082)	0.402 (.238)	0.140 (.241)	0.107 (.131)	0.126 (.051)	
Pred		0.284 (.622)	0.285 (.625)	0.250 (.513)	0.391 (.142)	0.918 (2.01)	0.921 (2.02)	0.631 (.976)	0.656 (.216)	
		$r = 2, p = 1000, \tau = 80\%$			$r = 5, p = 1000, \tau = 80\%$					
		als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
ase		1.250 (.700)	0.756 (.715)	0.590 (.644)	0.631 (.348)	2.161 (.570)	1.778 (.579)	1.200 (.280)	1.221 (.170)	
rse		0.624 (.334)	0.376 (.337)	0.291 (.292)	0.313 (.153)	0.434 (.123)	0.357 (.123)	0.241 (.060)	0.245 (.038)	
Pred		0.798 (.758)	0.799 (.760)	0.610 (.654)	0.598 (.281)	1.957 (.663)	1.958 (.664)	1.284 (.299)	1.272 (.160)	

The mean and the standard deviation (in parentheses) of each error between the simulation replicates are presented



Table 4 continued

		$r = 2, p = 1000, \tau = 20\%$			$r = 5, p = 1000, \tau = 20\%$					
		als	d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
		$r = 2, p = 1000, \tau = 50\%$			$r = 5, p = 1000, \tau = 50\%$					
	als		d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
ase	1.680 (.815)	0.244 (.842)	0.353 (1.98)	0.464 (.254)	1.799 (.337)	0.474 (.306)	0.419 (.121)	0.664 (.210)		
rse	0.110 (.036)	0.013 (.023)	0.016 (.054)	0.030 (.015)	0.045 (.011)	0.011 (.003)	0.010 (.004)	0.016 (.006)		
Pred	0.322 (1.38)	0.322 (1.39)	0.542 (3.66)	0.468 (.211)	0.613 (.600)	0.605 (.584)	0.501 (.186)	0.758 (.236)		
		$r = 2, p = 1000, \tau = 80\%$			$r = 5, p = 1000, \tau = 80\%$					
	als		d.b	f.Bayes	Bayes	als	d.b	f.Bayes	Bayes	
ase	6.350 (15.1)	5.880 (15.3)	6.414 (18.1)	0.849 (.346)	6.426 (10.1)	6.028 (10.2)	2.287 (2.96)	2.246 (2.20)		
rse	0.351 (.794)	0.321 (.802)	0.328 (.895)	0.051 (.015)	0.127 (.133)	0.117 (.135)	0.056 (.085)	0.050 (.017)		
Pred	7.188 (18.9)	7.198 (18.9)	7.887 (22.6)	0.876 (.403)	7.294 (12.6)	0.474 (.196)	2.646 (3.68)	2.528 (2.74)		

The mean and the standard deviation (in parentheses) of each error between the simulation replicates are presented



**Table 5** Simulation results for Setting I on empirical coverage rate of the confidence intervals and of the credible intervals of the entries (standard deviation is given in parentheses)

Missing rate	Empirical coverage rate (%)		
	CI db	CI f.Bayes	CI Bayes
<i>r</i> = 2, <i>p</i> = 100			
$\tau = 20\%$	99.9 (0.1)	99.9 (0.1)	99.3 (0.3)
$\tau = 50\%$	94.1 (1.1)	99.9 (0.1)	99.2 (0.3)
$\tau = 80\%$	58.9 (4.8)	98.6 (0.4)	98.5 (0.5)
<i>r</i> = 2, <i>p</i> = 1000			
$\tau = 20\%$	99.9 (0.1)	99.9 (0.1)	99.4 (0.1)
$\tau = 50\%$	94.5 (0.4)	99.9 (.01)	99.3 (0.1)
$\tau = 80\%$	64.4 (0.9)	99.6 (0.1)	99.0 (0.2)
<i>r</i> = 5, <i>p</i> = 100			
$\tau = 20\%$	99.9 (0.1)	99.9 (0.1)	99.2 (0.2)
$\tau = 50\%$	93.0 (1.3)	99.8 (0.1)	99.1 (0.2)
$\tau = 80\%$	38.5 (4.5)	96.4 (0.6)	90.9 (2.8)
<i>r</i> = 5, <i>p</i> = 1000			
$\tau = 20\%$	99.9 (0.1)	99.9 (0.1)	99.4 (0.1)
$\tau = 50\%$	94.1 (0.3)	99.9 (0.1)	99.2 (0.1)
$\tau = 80\%$	60.4 (0.8)	99.2 (0.1)	98.7 (0.1)

### 3.3 Results on uncertainty quantification

To examine the uncertainty quantification across the methods, we simulate a matrix as in Sect. 3.1 then we repeat the observation process 100 times. More precisely, we obtain 100 data sets by replicating the observation of 20%, 50% and 80% entries of the matrix  $M^*$  using a uniform sampling and then each sampled set is corrupted by noise as in (1), where the  $\mathcal{E}_i$  are i.i.d  $\mathcal{N}(0, 1)$ .

Tables 5 and 6 gather the empirical coverage rate of the confidence intervals and of the credible intervals of all methods over 100 independent experiments. More precisely, we report the 95% confidence intervals for the de-biased method. The credible interval is reported using the 89% equal-tailed interval, see [22, 32], for small posterior samples as in our situations with 500 posterior samples. We noted that to obtain 95% intervals, an effective posterior sample size of at least 10,000 is recommended [22]. A few examples from Setting I with 10000 posterior samples are given in Fig. 2.

A noteworthy conclusion from the results in Tables 5 and 6 is that the coverage rates of the 89% credible intervals are significantly higher than those of the 95% confidence intervals revealed by the de-biased method. The credible intervals of the 'f.Bayes' approach show a slightly better coverage rate than those based on the 'Bayes' approach. It is also noted that in the setting of approximate low-rankness, Table 6, where we do not have theoretical guarantee for the de-biased estimator, the coverage rate of the confidence intervals is very low while the credible intervals still come with reliable coverage rates. These results further explain why Bayesian methods yield better results in accuracy as in Tables 1 and 2.

In Fig. 2, we compare the limiting Gaussian distribution of the de-biased estimator and posterior samples for the 'f.Bayes' method against the true entries of interest. It is shown that the limiting Gaussian distribution of the de-biased estimator yields a slightly sharper tail distribution compared to the distribution of the posterior samples. In addition, displays the

**Table 6** Simulation results for Setting II on empirical coverage rate of the confidence intervals and of the credible intervals of the entries (standard deviation is given in parentheses)

Missing rate	Empirical coverage rate (%)		
	CI db	CI f.Bayes	CI Bayes
	Approximate rank-2, $p = 100$		
$\tau = 20\%$	70.4 (0.3)	69.3 (0.3)	55.9 (0.4)
$\tau = 50\%$	49.3 (0.4)	78.0 (0.4)	65.2 (0.5)
$\tau = 80\%$	33.0 (2.6)	89.5 (0.6)	81.3 (0.7)
	Approximate rank-2, $p = 1000$		
$\tau = 20\%$	57.2 (0.1)	56.0 (0.1)	51.1 (1.8)
$\tau = 50\%$	37.5 (1.3)	63.0 (0.2)	51.2 (1.1)
$\tau = 80\%$	27.6 (0.8)	77.6 (0.3)	67.3 (0.4)
	Approximate rank-5, $p = 100$		
$\tau = 20\%$	90.1 (0.3)	88.9 (0.3)	76.9 (0.4)
$\tau = 50\%$	66.9 (0.5)	93.2 (0.3)	83.8 (0.4)
$\tau = 80\%$	30.9 (3.1)	94.2 (0.6)	90.2 (2.8)
	Approximate rank-5, $p = 1000$		
$\tau = 20\%$	78.9 (0.1)	77.4 (0.1)	69.1 (1.1)
$\tau = 50\%$	56.3 (0.2)	85.0 (0.1)	73.4 (0.9)
$\tau = 80\%$	38.0 (0.7)	92.6 (0.2)	86.2 (0.3)

**Table 7** Simulation results for Setting III on empirical coverage rate of the confidence intervals and of the credible intervals of the entries (standard deviation is given in parentheses)

Missing rate	Empirical coverage rate (%)		
	CI db	CI f.Bayes	CI Bayes
	$r = 2, p = 100$		
$\tau = 20\%$	95.7 (5.2)	96.4 (4.8)	97.9 (0.4)
$\tau = 50\%$	71.4 (7.6)	94.9 (4.6)	98.2 (0.3)
$\tau = 80\%$	30.6 (5.8)	94.9 (2.0)	97.1 (0.8)
	$r = 2, p = 1000$		
$\tau = 20\%$	96.7 (4.5)	96.7 (3.4)	97.2 (0.4)
$\tau = 50\%$	75.6 (3.1)	95.9 (4.5)	97.7 (0.3)
$\tau = 80\%$	41.7 (4.0)	95.5 (2.0)	98.5 (0.2)
	$r = 5, p = 100$		
$\tau = 20\%$	95.9 (3.1)	96.4 (2.5)	95.2 (0.8)
$\tau = 50\%$	70.5 (3.2)	96.5 (1.6)	95.7 (0.7)
$\tau = 80\%$	24.6 (2.3)	95.3 (0.8)	94.0 (1.6)
	$r = 5, p = 1000$		
$\tau = 20\%$	96.9 (1.7)	96.8 (2.5)	93.9 (0.6)
$\tau = 50\%$	74.9 (1.8)	96.6 (1.6)	94.6 (0.4)
$\tau = 80\%$	39.6 (2.1)	96.1 (0.6)	95.9 (0.5)

**Table 8** Simulation results for Setting IV on empirical coverage rate of the confidence intervals and of the credible intervals of the entries (standard deviation is given in parentheses)

Missing rate	Empirical coverage rate (%)		
	CI db	CI f.Bayes	CI Bayes
$r = 2, p = 100$			
$\tau = 20\%$	97.1 (1.1)	96.4 (4.7)	97.9 (0.5)
$\tau = 50\%$	75.4 (4.9)	95.6 (6.3)	98.1 (0.3)
$\tau = 80\%$	33.4 (7.1)	94.7 (2.8)	97.4 (0.7)
$r = 2, p = 1000$			
$\tau = 20\%$	97.0 (0.6)	97.2 (0.4)	97.2 (0.4)
$\tau = 50\%$	76.0 (4.1)	97.2 (0.4)	97.9 (0.3)
$\tau = 80\%$	39.1 (6.0)	93.8 (5.9)	98.6 (0.2)
$r = 5, p = 100$			
$\tau = 20\%$	97.0 (0.9)	96.9 (2.2)	95.7 (0.7)
$\tau = 50\%$	72.4 (2.7)	96.7 (0.9)	96.0 (0.6)
$\tau = 80\%$	22.9 (3.3)	92.9 (1.4)	94.5 (1.0)
$r = 5, p = 1000$			
$\tau = 20\%$	97.1 (0.3)	97.2 (0.2)	94.6 (0.5)
$\tau = 50\%$	75.3 (0.7)	96.9 (1.2)	94.8 (0.4)
$\tau = 80\%$	37.3 (2.6)	96.2 (0.5)	96.3 (0.3)

Q-Q (quantile-quantile) plots of 10,000 posterior samples of some entries vs. the standard Gaussian random variables. It shows that the posterior distributions of these entries reasonably well match the standard Gaussian distribution.

Results on empirical coverage rate of the confidence intervals and of the credible intervals for Setting III and IV with heavy tail noises are gathered in Tables 7 and 8. We can see that there is a slight reduction in the empirical coverage rate of all methods compared with those in the Gaussian noise setting in Table 5. As in Setting I and II, the empirical coverage rates of confidence intervals decrease quickly as the missing rates  $\tau$  increase, while the empirical coverage rates of credible intervals remain stable.

## 4 Discussion and conclusion

In this paper, we have provided extensive numerical comparisons between the de-biased estimator and the Bayesian estimators in the problem of low-rank matrix completion. Results from numerical simulations draw a systematic picture of the behaviour of these estimators originally. More specifically, on the estimation accuracy, the de-biased estimator is comparable to the Bayesian estimators whereas the Bayesian estimators are much more stable and in some cases can outperform the de-biased estimator, especially in the small samples regime. Moreover, the credible intervals reasonably cover the underlying entries quite well and slightly better than the confidence intervals in exact low-rank matrix completion. However, in the case of approximate low-rankness, the confidence intervals revealed by the de-biased estimators no longer work well. These results are interested for and can be served as a guideline for researchers as well as practitioners in many areas where one only has access to a few observations.

On the other hand, the results in this work suggest that the considered Bayesian estimators may actually reach the minimax-optimal rate of convergence without additional logarithmic factor. The extra log-terms could be due to the PAC-Bayesian bounds technique that used to prove the theoretical properties of the Bayesian estimator. Moreover, as shown in Alquier and Ridgway [3], the same rate with log-term is proved for the concentration of the corresponding posterior and we conjecture that this rate could also be improved due to the coverage of credible intervals. These are important questions that remain open up to our knowledge.

Last but not least, it is also important to perform the comparisons with the Variational Bayesian (VB) method in Lim and Teh [24] where its theoretical guarantees are given in Lim and Teh [3], because this method is very popular for matrix completion with large datasets. This will be the objective of our future work. However, we would like to note that, in a preprint [4], the authors had performed some comparisons between the Bayesian approach and the VB method. The message from their works is that we can expect that VB should be more or less as accurate as Bayes, maybe slightly less, but that the credibility intervals would be inaccurate (see e.g Figure 3 in Alquier et al. [4]).

**Acknowledgements** The author would like to thank the editor and the anonymous referee, who kindly reviewed the earlier version of the manuscript, for providing valuable suggestions and enlightening comments that help improve the current version of the paper. TTM is supported by the Norwegian Research Council grant number 309960 through the Centre for Geophysical Forecasting at NTNU. The author would like to thank Pierre Alquier for kindly providing useful feedbacks on a first draft of this paper.

**Funding** Open access funding provided by NTNU Norwegian University of Science and Technology (incl St. Olavs Hospital - Trondheim University Hospital)

**Data availability** The R codes and data used in the numerical experiments are available at: <https://github.com/tienmt/UQMC>.

## Declarations

**Conflict of interest** The authors declare that they have no conflict of interest.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Adomavicius, G., Tuzhilin, A.: Context-aware recommender systems. In: Recommender Systems Handbook, pp. 217–253. Springer, Berlin (2011)
2. Alquier, P.: Bayesian methods for low-rank matrix estimation: short survey and theoretical study. In: Algorithmic Learning Theory 2013. Springer, Berlin, pp 309–323 (2013)
3. Alquier, P., Ridgway, J.: Concentration of tempered posteriors and of their variational approximations. *Ann. Stat.* **48**(3), 1475–1497 (2020)
4. Alquier, P., Cottet, V., Chopin, N., Rousseau, J.: Bayesian matrix completion: prior specification and consistency. (2014) arXiv preprint [arXiv:1406.1440](https://arxiv.org/abs/1406.1440)
5. Babacan, S.D., Luessi, M., Molina, R., Katsaggelos, A.K.: Sparse Bayesian methods for low-rank matrix estimation. *IEEE Trans. Signal Process.* **60**(8), 3964–3977 (2012)

6. Bennett, J., Lanning, S.: The netflix prize. In: Proceedings of KDD Cup and Workshop, vol 2007, p 35 (2007)
7. Cabral, R., De la Torre, F., Costeira, J.P., Bernardino, A.: Matrix completion for weakly-supervised multi-label image classification. *IEEE Trans. Pattern Anal. Mach. Intell.* **37**(1), 121–135 (2014)
8. Candès, E.J., Plan, Y.: Matrix completion with noise. *Proc. IEEE* **98**(6), 925–936 (2010)
9. Candès, E.J., Recht, B.: Exact matrix completion via convex optimization. *Found. Comput. Math.* **9**(6), 717–772 (2009)
10. Candès, E.J., Tao, T.: The power of convex relaxation: near-optimal matrix completion. *IEEE Trans. Inform. Theory* **56**(5), 2053–2080 (2010)
11. Chen, Y., Fan, J., Ma, C., Yan, Y.: Inference and uncertainty quantification for noisy matrix completion. *Proc. Natl. Acad. Sci.* **116**(46), 22931–22937 (2019)
12. Chi, E.C., Zhou, H., Chen, G.K., Del Vecchio, D.O., Lange, K.: Genotype imputation via matrix completion. *Genome Res.* **23**(3), 509–518 (2013)
13. Cottet, V., Alquier, P.: 1-bit matrix completion: Pac-bayesian analysis of a variational approximation. *Mach. Learn.* **107**(3), 579–603 (2018)
14. Fithian, W., Mazumder, R.: Flexible low-rank statistical modeling with missing data and side information. *Stat. Sci.* **33**(2), 238–260 (2018)
15. Gross, D., Liu, Y.-K., Flammia, S.T., Becker, S., Eisert, J.: Quantum state tomography via compressed sensing. *Phys. Rev. Lett.* **105**(15), 150401 (2010)
16. Harper, F.M., Konstan, J.A.: The movielens datasets: history and context. *Acm Trans. Interact. Intell. Syst. (TIIS)* **5**(4), 1–19 (2015)
17. Hastie, T., Mazumder, R., Lee, J.D., Zadeh, R.: Matrix completion and low-rank SVD via fast alternating least squares. *J. Mach. Learn. Res.* **16**(1), 3367–3402 (2015)
18. He, K., Sun, J.: Image completion approaches using the statistics of similar patches. *IEEE Trans. Pattern Anal. Mach. Intell.* **36**(12), 2423–2435 (2014)
19. Jiang, B., Ma, S., Causey, J., Qiao, L., Hardin, M.P., Bitts, I., Johnson, D., Zhang, S., Huang, X.: Sparrec: an effective matrix completion framework of missing data imputation for GWAS. *Sci. Rep.* **6**(1), 1–15 (2016)
20. Klopp, O.: Matrix completion by singular value thresholding: sharp bounds. *Electron. J. Stat.* **9**(2), 2348–2369 (2015)
21. Koltchinskii, V., Lounici, K., Tsybakov, A.B.: Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Stat.* **39**(5), 2302–2329 (2011)
22. Kruschke, J.: *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan.* Academic Press, New York (2014)
23. Lawrence, N.D., Urtasun, R.: Non-linear matrix factorization with gaussian processes. In: Proceedings of the 26th Annual International Conference on Machine Learning. ACM, pp. 601–608 (2009)
24. Lim, Y.J., Teh, Y.W.: Variational bayesian approach to movie rating prediction. In: Proceedings of KDD Cup and Workshop, vol. 7, pp. 15–21 (2007)
25. Luo, Y., Liu, T., Tao, D., Xu, C.: Multiview matrix completion for multilabel image classification. *IEEE Trans. Image Process.* **24**(8), 2355–2368 (2015)
26. Mai, T.T.: Bayesian matrix completion with a spectral scaled student prior: theoretical guarantee and efficient sampling. (2021a). arXiv preprint [arXiv:2104.08191](https://arxiv.org/abs/2104.08191)
27. Mai, T.T.: dbMC: Confidence Interval for Matrix Completion via De-Biased Estimator. R package version 1.0.0 (2021b)
28. Mai, T.T.: An efficient adaptive mcmc algorithm for pseudo-bayesian quantum tomography. *Comput. Stat.*, 1–17 (2022)
29. Mai, T.T., Alquier, P.: A Bayesian approach for noisy matrix completion: optimal rate under general sampling distribution. *Electron. J. Stat.* **9**(1), 823–841 (2015)
30. Mai, T.T., Alquier, P.: Pseudo-Bayesian quantum tomography with rank-adaptation. *J. Stat. Plan. Infer.* **184**, 62–76 (2017)
31. Mazumder, R., Hastie, T., Tibshirani, R.: Spectral regularization algorithms for learning large incomplete matrices. *J. Mach. Learn. Res.* **11**, 2287–2322 (2010)
32. McElreath, R.: *Statistical Rethinking: A Bayesian Course with Examples in R and Stan.* CRC press, Boca Raton (2020)
33. Negahban, S., Wainwright, M.J.: Restricted strong convexity and weighted matrix completion: optimal bounds with noise. *J. Mach. Learn. Res.* **13**, 1665–1697 (2012)
34. Recht, B., Ré, C.: Parallel stochastic gradient algorithms for large-scale matrix completion. *Math. Program. Comput.* **5**(2), 201–226 (2013)
35. Rendle, S., Zhang, L., Koren, Y.: On the difficulty of evaluating baselines: A study on recommender systems (2019). arXiv preprint [arXiv:1905.01395](https://arxiv.org/abs/1905.01395)

36. Rendle, S., Krichene, W., Zhang, L., Anderson, J.: Neural collaborative filtering vs. matrix factorization revisited. In: Fourteenth ACM Conference on Recommender Systems, pp. 240–248 (2020)
37. Salakhutdinov, R., Mnih, A.: Bayesian probabilistic matrix factorization using markov chain monte carlo. In: Proceedings of the 25th International Conference on Machine learning. ACM, pp. 880–887 (2008)
38. Xiong, L., Chen, X., Huang, T.-K., Schneider, J., Carbonell, J.G.: Temporal collaborative filtering with bayesian probabilistic tensor factorization. In: Proceedings of the 2010 SIAM International Conference on Data Mining. SIAM, pp. 211–222 (2010)
39. Yang, L., Fang, J., Duan, H., Li, H., Zeng, B.: Fast low-rank Bayesian matrix completion with hierarchical gaussian prior models. *IEEE Trans. Signal Process.* **66**(11), 2804–2817 (2018)
40. Zhou, M., Wang, C., Chen, M., Paisley, J., Dunson, D., Carin, L.: Nonparametric bayesian matrix completion. In: 2010 IEEE Sensor Array and Multichannel Signal Processing Workshop. IEEE, pp. 213–216 (2010)

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.