
Private Alternating Least Squares: Practical Private Matrix Completion with Tighter Rates

Steve Chien¹ Prateek Jain¹ Walid Krichene^{*1} Steffen Rendle¹
Shuang Song¹ Abhradeep Thakurta^{*1} Li Zhang¹

Abstract

We study the problem of differentially private (DP) matrix completion under user-level privacy. We design a joint differentially private variant of the popular Alternating-Least-Squares (ALS) method that achieves: i) (nearly) optimal sample complexity for matrix completion (in terms of number of items, users), and ii) the best known privacy/utility trade-off both theoretically, as well as on benchmark data sets. In particular, we provide the first global convergence analysis of ALS with *noise* introduced to ensure DP, and show that, in comparison to the best known alternative (the Private Frank-Wolfe algorithm by Jain et al. (2018)), our error bounds scale significantly better with respect to the number of items and users, which is critical in practical problems. Extensive validation on standard benchmarks demonstrate that the algorithm, in combination with carefully designed sampling procedures, is significantly more accurate than existing techniques, thus promising to be the first practical DP embedding model.

1. Introduction

Given $M_{ij}, (i, j) \in \Omega$ where $\Omega \subseteq [n] \times [m]$ is a set of observed user-item ratings, and assuming $M \approx U^*(V^*)^\top \in \mathbb{R}^{n \times m}$ to be a nearly low-rank matrix, the goal of low-rank matrix completion (LRMC) is to efficiently learn $\hat{U} \in \mathbb{R}^{n \times r}$ and $\hat{V} \in \mathbb{R}^{m \times r}$, such that $M \approx \hat{U}\hat{V}^\top$.

LRMC, a.k.a. matrix factorization, is a cornerstone technique for building recommendation systems (Koren & Bell, 2015; Hu et al., 2008), and though proposed over a decade ago, it remains highly competitive (Rendle et al., 2019). In the recommendation setting, M represents a mostly un-

known user-item ratings matrix and \hat{U} and \hat{V} capture the user and item embeddings. Using the learned (\hat{U}, \hat{V}) , the system computes rating predictions $\hat{M}_{ij} = (\hat{U}\hat{V}^\top)_{ij}$ to recommend items for the users. To ensure good generalization, one would set the rank $r \ll \min(m, n)$.

Such models, while highly successful in practice, have the risk of leaking users' ratings through model parameters or their recommendations. The privacy risk of similar models has been well documented, and the protection against it has been intensively studied (Dinur & Nissim, 2003; Dwork et al., 2007; Korolova, 2010; Calandrino et al., 2011; Shokri et al., 2017; Carlini et al., 2019; 2020a;b; Thakkar et al., 2020). In this paper, we focus on learning user and item embeddings, and consequently user-item recommendations, while ensuring privacy of users' ratings.

We conform to the well-established formal notion of differential privacy (DP) (Dwork et al., 2006a;b) to protect users' ratings. We operate in the setting of *user-level* privacy (Dwork & Roth, 2014; Jain et al., 2018), where we intend to protect *all the ratings by the user*, a much harder task than protecting a single rating from the user (a.k.a. *entry-level privacy*) (Hardt & Roth, 2013; Meng et al., 2018). Note that *user-level* privacy is critical in this problem, as the ratings from a single user tend to be correlated and can thus be used to fingerprint a user (Calandrino et al., 2011). As is standard in the user-level privacy literature (Jain et al., 2018), we estimate the shared item embeddings \hat{V} while preserving privacy with respect to the users. In contrast, each user *independently* computes their embedding (a row of \hat{U}) as a function of their own ratings and the privacy preserving item embeddings \hat{V} . Formally, this setup is called *joint differential privacy* (Kearns et al., 2014), and it is well-established (Hardt & Roth, 2012; 2013) that such a relaxation is necessary to learn non-trivial recommendations while ensuring user-level privacy.

While several works have studied LRMC under joint-differential privacy (McSherry & Mironov, 2009; Liu et al., 2015; Jain et al., 2018), most of the existing techniques do not provide satisfactory empirical performance compared to the state-of-the-art (SOTA) non-private LRMC methods. Furthermore, these works either lack a rigorous performance

¹Google Research. Correspondence to: Walid Krichene <walidk@google.com>, Abhradeep Thakurta <athakurta@google.com>.

analysis (McSherry & Mironov, 2009; Liu et al., 2015) or provide guarantees that are significantly weaker (Jain et al., 2018) than that of non-private LRMC algorithms. Matrix factorization can also be solved using other first-order methods such as stochastic gradient descent (Ge et al., 2016) or alternating gradient descent (Lu et al., 2019), so one may apply the differentially private SGD (DPSGD) algorithm (Song et al., 2013; Bassily et al., 2014; Abadi et al., 2016) to achieve privacy. However, applying DPSGD to LRMC is challenging as SGD typically requires many steps to converge, thus increasing privacy cost.

In this work, we design and analyze a differentially private version of the widely used alternating least squares (ALS) algorithm for LRMC (Koren et al., 2009; Jain et al., 2013). ALS alternates between optimizing over the user embeddings \hat{U} and the item embeddings \hat{V} , each through least squares minimization. One important property of ALS is that when solving for one side, the optimization can be done independently for each user or item, which makes ALS highly scalable. Our key insight is that this decoupling of the solution is also useful for privacy-preserving computation, since there is no accumulation of noise when solving for the embeddings of different users (or items). Besides, ALS is known to require few iterations to converge in practice, making it particularly suitable for privacy preserving LRMC.

Indeed, we present a differentially private variant of ALS, which we refer to as DPALS, and demonstrate that it enjoys much tighter error rates (see Table 1) and better empirical performance than the current SOTA, the differentially private Frank-Wolfe (DPFW) method of Jain et al. (2018). Furthermore, on the large scale benchmark of MovieLens 20M, DPALS produces the first realistic DP embedding model with competitive recall metric under moderate privacy loss.

More specifically, our contributions are the following.

Private alternating least squares for matrix completion.

We provide the first differentially private version of alternating least squares (DPALS) for matrix completion with user-level privacy guarantee (Section 3). The algorithm is conceptually simple, efficient, and highly scalable. We provide rigorous analysis on its privacy guarantee under the notion of Joint Rényi Differential Privacy.

Initialization via noisy power iteration. For convergence of DPALS algorithm, we need it to be initialized with a \hat{V}^0 close to V^* in spectral norm. The standard approach based on private PCA (Dwork et al., 2014) would require $n = \tilde{\Omega}\left(\frac{m\sqrt{m}}{\varepsilon}\right)$ to achieve the initialization condition. Instead, we show that with a careful analysis, initializing with noisy power iteration only requires $n = \tilde{\Omega}\left(\frac{m}{\varepsilon}\right)$. Our analysis shows in particular that it suffices that the top- r eigenspace of $A := P_{\Omega}(\mathbf{M})^{\top} P_{\Omega}(\mathbf{M})$ be incoherent, and that there be a $\Omega(\log^2 m)$ gap between the top- r eigenvalues

Table 1. Sample complexity bounds for various algorithms, assuming constant Frobenius norm error. Here, n is the number of users, m is the number of items, and $\tilde{\Omega}(\cdot)$ hides $\text{polylog}(n, m, 1/\delta)$. (*) assumes additional property of \mathbf{M} being incoherent.

Algorithm	Bound on n	Bound on $ \Omega /n$	Iterations
Trace Norm (*) (non-priv.) (Candès & Recht, 2009)	$\tilde{\Omega}(m)$	$\tilde{\Omega}(\log^2 n)$	$\text{poly}(n, m)$
ALS (*) (non-priv.) (Jain et al., 2013)	$\tilde{\Omega}(m)$	$\tilde{\Omega}(\log^2 n)$	$\text{polylog}(n, m)$
Private SVD(*) (McSherry & Mironov, 2009)	-	-	-
Private SGLD (Liu et al., 2015)	-	-	-
Private FW (Jain et al., 2018)	$\tilde{\Omega}(m^{5/4})$	$\tilde{\Omega}(\sqrt{m})$	$\text{poly}(n, m)$
Private ALS (*) (this work)	$\tilde{\Omega}(m)$	$\tilde{\Omega}(\log^3 n)$	$\text{polylog}(n, m)$

and the rest. This result improves on (Hardt & Price, 2013) which required all the eigenvectors of \mathbf{A} to be incoherent, a condition that is hard to guarantee in our setting (Dekel et al., 2011; Vu & Wang, 2015; Rudelson & Vershynin, 2015).

Tighter privacy/utility/computation trade-offs. We prove theoretical guarantees on the sample complexity and the error bounds of DPALS under standard assumptions (Section 4). These bounds are much tighter than the current SOTA, the DPFW method (Jain et al., 2018). In particular, we show the following. First, DPALS requires only $O(\log^{O(1)} n)$ samples per user to guarantee its convergence. In contrast, DPFW requires \sqrt{m} ratings per user. Second, to achieve a Frobenius norm error of ζ , DPALS requires $n = \tilde{\Omega}\left(\frac{m\sqrt{m}}{\zeta\varepsilon} + m\right)$ users, which is nearly optimal in terms of ζ and ε . In contrast, DPFW’s sample complexity is $n = \tilde{\Omega}\left(m^{5/4}/(\zeta^5\varepsilon)\right)$; note a significant improvement in terms of ζ . Finally, Private SVD (McSherry & Mironov, 2009) is not even *consistent*, i.e., for a fixed $\varepsilon, m, |\Omega| = n\sqrt{m}$, even if we scale $n \rightarrow \infty$, the Frobenius norm error bound does not converge to 0 (see Theorem B.3 of Jain et al. (2018)).

Practical techniques to improve accuracy. One main difficulty in applying DPALS to practical problems comes from a heavy skew in the item distribution. We propose two heuristics to reduce the skew while preserving privacy (Section 5). Experiments on real-world benchmarks show that these techniques can significantly improve model quality.

Strong empirical results using DPALS. We carry out an extensive study of DPALS on synthetic and real-world benchmarks. Aided by the aforementioned practical techniques, DPALS achieves significant gains over the current SOTA method. In particular, on the MovieLens 10M rating prediction benchmark, DPALS achieves the same error rate as the current SOTA even when trained on a fraction (23%) of users. When trained on all users, it achieves a relative decrease in RMSE of at least 7%. DPALS also achieves remarkably good performance on the MovieLens 20M item recommendation benchmark with modest privacy loss, and remains competitive even with non-private ALS, the first DP private embedding model to achieve such strong results.

2. Background

2.1. Notation

Let $[m]$ denote the set $\{1, 2, \dots, m\}$. Let $\mathbb{R}^{n \times m}$ denote the set of $n \times m$ matrices. Throughout the paper, we use bold face uppercase letters to represent matrices and lowercase letters for vectors. For any matrix $\mathbf{A} = (\mathbf{A}_{ij}) \in \mathbb{R}^{n \times m}$, let \mathbf{A}_i be the i -th row vector of \mathbf{A} . Denote by $\|\mathbf{A}\|_F, \|\mathbf{A}\|_\infty$ the Frobenius norm and the max norm of \mathbf{A} . For $\Omega \subseteq [n] \times [m]$, define the projection $\mathbf{P}_\Omega(\mathbf{A}) \in \mathbb{R}^{n \times m}$ as $\mathbf{P}_\Omega(\mathbf{A})_{ij} = \mathbf{A}_{ij}$ if $(i, j) \in \Omega$ and 0 otherwise. For $i \in [n]$, define $\Omega_i := \{j : (i, j) \in \Omega\}$. Similarly, for $j \in [m]$, let $\Omega_j = \{i : (i, j) \in \Omega\}$. For $\mathbf{u}, \mathbf{v} \in \mathbb{R}^r$, we use $\mathbf{u} \cdot \mathbf{v} \in \mathbb{R}$ to denote their dot product, and $\mathbf{u} \otimes \mathbf{v} \in \mathbb{R}^{r \times r}$ for their outer product.

2.2. Matrix Completion, Alternating Least Squares

Let $\mathbf{M} \in \mathbb{R}^{n \times m}$ be a rank r matrix, such that each entry M_{ij} ($i \in [n], j \in [m]$) represents the preference/affinity of user i for item j . Given a set of observed entries $\mathbf{P}_\Omega(\mathbf{M})$, $\Omega \subseteq [n] \times [m]$, the goal of LRMC is to reconstruct \mathbf{M} with minimal error. This can be achieved by finding $\hat{\mathbf{U}} \in \mathbb{R}^{n \times r}$ and $\hat{\mathbf{V}} \in \mathbb{R}^{m \times r}$ such that the regularized squared error $\|\mathbf{P}_\Omega(\mathbf{M} - \hat{\mathbf{U}}\hat{\mathbf{V}}^\top)\|_F^2 + \lambda\|\hat{\mathbf{U}}\|_F^2 + \lambda\|\hat{\mathbf{V}}\|_F^2$ is minimized. This minimization problem is NP-hard in general (Hardt et al., 2014). But the alternating least squares (ALS) algorithm has proved to work well in practice.

ALS alternately computes $\hat{\mathbf{U}}, \hat{\mathbf{V}}$ by minimizing the above objective while assuming the other embeddings fixed. Each step can be solved efficiently through the standard least squares algorithm with the following closed form solution.

$$\forall i \quad \hat{\mathbf{U}}_i^t = (\lambda \mathbf{I} + \sum_{j \in \Omega_i} \hat{\mathbf{V}}_j^t \otimes \hat{\mathbf{V}}_j^t)^{-1} \sum_{j \in \Omega_i} M_{ij} \hat{\mathbf{V}}_j^t, \quad (1)$$

$$\forall j \quad \hat{\mathbf{V}}_j^{t+1} = (\lambda \mathbf{I} + \sum_{i \in \Omega_j} \hat{\mathbf{U}}_i^t \otimes \hat{\mathbf{U}}_i^t)^{-1} \sum_{i \in \Omega_j} M_{ij} \hat{\mathbf{U}}_i^t. \quad (2)$$

While ALS does not guarantee convergence to the global optimum in general, it works remarkably well in practice and often produces $\hat{\mathbf{U}}$ and $\hat{\mathbf{V}}$ such that $\hat{\mathbf{U}}\hat{\mathbf{V}}^\top$ is a good approximation of \mathbf{M} . The practical success of ALS has inspired many theoretical analyses, which make the following additional assumptions on \mathbf{M} and Ω .

Assumption 1 (μ -incoherence). Let $\mathbf{M} = \mathbf{U}^* \mathbf{\Sigma}^* (\mathbf{V}^*)^\top$ be the singular value decomposition of \mathbf{M} , i.e. $\mathbf{U}^* \in \mathbb{R}^{n \times r}$, $\mathbf{V}^* \in \mathbb{R}^{m \times r}$ are orthonormal matrices, and $\mathbf{\Sigma}^* \in \mathbb{R}^{r \times r}$ is the diagonal matrix of the singular values of \mathbf{M} . We assume that \mathbf{M} is μ -incoherent, that is, $\forall i \in [n]$, $\|\mathbf{U}_i^*\|_2 \leq \frac{\mu\sqrt{r}}{\sqrt{n}}$; and $\forall j \in [m]$, $\|\mathbf{V}_j^*\|_2 \leq \frac{\mu\sqrt{r}}{\sqrt{m}}$.

Assumption 2 (Random Ω). We assume that Ω are random observations with probability p , that is, $\Omega = \{(i, j) \in [n] \times [m] : \delta_{ij} = 1\}$, where $\delta_{ij} \in \{0, 1\}$ are i.i.d. random variables with $\Pr[\delta_{ij} = 1] = p$.

Jain et al. (2013); Hardt & Wooters (2014) showed that ALS converges to \mathbf{M} with high probability if \mathbf{M} is μ -incoherent and $p = \tilde{\Omega} \left(\frac{\log n}{m} \right)$, where $n \geq m$ and $\tilde{\Omega}$ hides polynomial dependence on μ, r , and the condition number of \mathbf{M} . In this work, we make the same assumptions on \mathbf{M} and Ω . Our key theoretical contribution is a similar convergence result for DPALS, under the additional requirements of user-level differential privacy.

2.3. Joint Differential Privacy

Differential privacy (Dwork et al., 2006b;a) is a widely adopted privacy notion. We use the variant of *user-level* joint differential privacy (Joint DP). Intuitively, Joint DP requires any information which may cross different users to be differentially private, but allows each individual user to use her own private information to her full advantage, for example, when computing the embeddings for generating recommendations to herself. This notion was already implicit in (McSherry & Mironov, 2009) and made formal in (Kearns et al., 2014; Jain et al., 2018).

Let $D = \{d_1, \dots, d_n\}$ be a data set of n records, where each sample d_i is drawn from a domain τ and belongs to individual i (which we also refer to as a *user*). Let $\mathcal{A} : \tau^* \rightarrow \mathcal{S}^n$ be an algorithm that produces n outputs in some space \mathcal{S} , one for each user i . Let D_{-i} be the data set with the i -th user removed, and let $\mathcal{A}_{-i}(D)$ be the set of outputs without that of the i -th user. Also, let $(d_i; D_{-i})$ be the data set obtained by adding d_i (for user i) to the data set D_{-i} . Joint DP and its Rényi differential privacy (Mironov, 2017) (Joint RDP) variant are defined as follows.

Definition 3 (Joint Differential Privacy (Kearns et al., 2014)). An Algorithm \mathcal{A} is (ϵ, δ) -jointly differentially private if for any user i , for any possible value of data entry $d_i, d'_i \in \tau$, for any instantiation of the data set for other users $D_{-i} \in \tau^{n-1}$, and for any set of outputs $S \subseteq \mathcal{S}^n$, the following two inequalities hold simultaneously:

$$\Pr_{\mathcal{A}}[\mathcal{A}_{-i}((d_i; D_{-i})) \in S] \leq e^\epsilon \Pr_{\mathcal{A}}[\mathcal{A}_{-i}(D_{-i}) \in S] + \delta$$

$$\Pr_{\mathcal{A}}[\mathcal{A}_{-i}(D_{-i}) \in S] \leq e^\epsilon \Pr_{\mathcal{A}}[\mathcal{A}_{-i}((d_i; D_{-i})) \in S] + \delta.$$

An algorithm \mathcal{A} is (α, ϵ) -joint Rényi differentially private (Joint RDP) if $D_\alpha(\mathcal{A}_{-i}((d_i; D_{-i})) \parallel \mathcal{A}_{-i}(D_{-i})) \leq \epsilon$ and $D_\alpha(\mathcal{A}_{-i}(D_{-i}) \parallel \mathcal{A}_{-i}((d_i; D_{-i}))) \leq \epsilon$, where D_α is the Rényi divergence of order α .

If we replace \mathcal{A}_{-i} with \mathcal{A} in the definition, we would recover the standard definition of DP and RDP. We note that the joint DP (resp. joint RDP) enjoys the same composability properties as the notion of DP (resp. RDP).

3. DPALS: Private Alternating Least Squares

We now provide the details of the DPALS algorithm and prove its privacy guarantee in the joint DP model.

Algorithm 1 DPALS: Private Matrix Completion via Alternating Minimization

Required: $P_\Omega(M)$: observed ratings, σ : noise standard deviation, Γ_u : row clipping parameter, Γ_M : entry clipping parameter, T : number of steps, λ : regularization parameter, r : rank, k : maximum number of ratings per user in $\mathcal{A}_{\text{item}}$, \hat{V}^0 : initial \hat{V} .

```

1 Clip entries in  $P_\Omega(M)$  so that  $\|P_\Omega(M)\|_\infty \leq \Gamma_M$ 
  for  $0 \leq t \leq T$  do
    for  $1 \leq i \leq n$  do
      2  $\hat{U}_i^t \leftarrow \mathcal{A}_{\text{user}}(\hat{V}^t, \Omega_i, P_\Omega(M)_i, T, \lambda, \Gamma_u)$ 
    end
    3  $\hat{U}^t \leftarrow [\hat{U}_1^t, \dots, \hat{U}_n^t]^\top$ 
    4  $\hat{V}^{t+1} \leftarrow \mathcal{A}_{\text{item}}(\hat{U}^t, \Omega, P_\Omega(M), k, \lambda, \Gamma_u, \Gamma_M)$ 
  end
5 return  $\hat{U}^T, \hat{V}^T$ 
    
```

Procedure $\mathcal{A}_{\text{item}}(U, \Omega, P_\Omega(M), k, \lambda, \Gamma_u, \Gamma_M)$

```

6  $\Omega' \leftarrow$  up to  $k$  random samples of  $(i, j) \in \Omega, \forall i \in [n]$ .
  for  $1 \leq j \leq m$  do
    7  $G_j \leftarrow \mathcal{N}_{\text{sym}}(0, \Gamma_u^4 \cdot \sigma^2)^{r \times r}$ 
    8  $g_j \leftarrow \mathcal{N}(0, \Gamma_u^2 \Gamma_M^2 \cdot \sigma^2)^r$ 
    9  $X_j \leftarrow \lambda I + \sum_{i \in \Omega'_j} U_i \otimes U_i + G_j$ 
    10  $V_j \leftarrow \Pi_{\text{PSD}}(X_j)^+ \left( \sum_{i \in \Omega'_j} M_{ij} \cdot U_i + g_j \right)$ 
  end
11  $\tilde{V} = [V_1, \dots, V_m]^\top$ 
12 return  $V = \tilde{V}(\tilde{V}^\top \tilde{V})^{-1/2}$ 
    
```

Procedure $\mathcal{A}_{\text{user}}(V, \Omega_i, P_\Omega(M)_i, T, \lambda, \Gamma_u)$

```

13  $\Omega'_i \leftarrow$  random samples of  $1/T$  fraction of  $j \in \Omega_i$ 
14  $u \leftarrow (\lambda I + \sum_{j \in \Omega'_i} V_j \otimes V_j)^{-1} \sum_{j \in \Omega'_i} M_{ij} V_j$ 
15 return clip( $u, \Gamma_u$ )
    
```

Notation. Let $\mathcal{N}(0, \sigma^2)$ be the Gaussian distribution of variance σ^2 , and $\mathcal{N}_{\text{sym}}(0, \sigma^2)^{r \times r}$ be the distribution of symmetric matrices where each entry in the upper triangle is drawn i.i.d. from $\mathcal{N}(0, \sigma^2)$. For a symmetric A , let $\Pi_{\text{PSD}}(A)$ be its projection to the positive semi-definite cone, obtained by replacing its negative eigenvalues with 0. Define $\text{clip}(u, c) = u \cdot \max(1, c/\|u\|_2)$, i.e., the projection of u on an ℓ_2 ball of radius c . Let A^\dagger be the pseudoinverse of A .

3.1. Algorithm

The private alternating least squares algorithm, DPALS, is described in Algorithm 1. It follows the standard ALS steps, i.e. it alternately solves the least squares problem to obtain \hat{U} and \hat{V} using (1) and (2). To guarantee joint DP, we compute differentially private item embeddings \hat{V}^{t+1} (using procedure $\mathcal{A}_{\text{item}}$) by solving a private variant of (2),

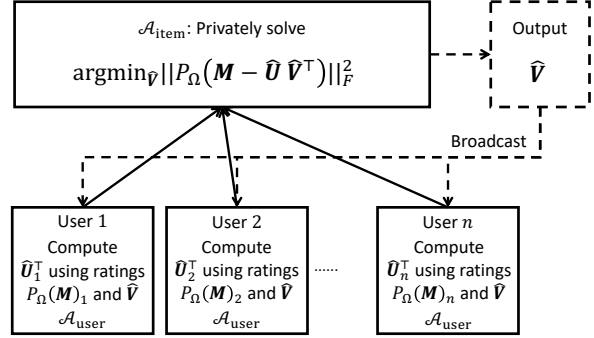


Figure 1. Block schematic of Joint differentially private alternating least squares algorithm. Solid lines and boxes represent privileged computations not visible to an adversary or other users. Dashed boxes and lines are public information accessible to anyone.

and compute each row of \hat{U}^{t+1} *independently without any noise* using procedure $\mathcal{A}_{\text{user}}$. A block schematic of the algorithm is presented in Figure 1.

Here we describe how the privacy is guaranteed in $\mathcal{A}_{\text{item}}$; see Theorem 1 for a formal statement. For a given $j \in [m]$, write $H_j^t = \lambda I + \sum_{i \in \Omega_j} \hat{U}_i^t \otimes \hat{U}_i^t$ and $w_j^t = \sum_{i \in \Omega_j} M_{ij} \hat{U}_i^t$. Then the non-private update (2) can be written as $\hat{V}_j^{t+1} = (H_j^t)^{-1} w_j^t$. In the private version, we need to add noise to protect both H_j^t and w_j^t . To ensure sufficient noise, we limit the influence of each user by “clipping” each \hat{U}_i^t to a bounded ℓ_2 norm Γ_u (Line 15 in $\mathcal{A}_{\text{user}}$) and resampling Ω such that each user participates in at most k items’ computation (Line 6 in $\mathcal{A}_{\text{item}}$). We then apply the Gaussian mechanism to H_j^t and w_j^t before using them to compute \hat{V}_j^{t+1} in $\mathcal{A}_{\text{item}}$ (Lines 9–12).

While the above procedure is sufficient to guarantee privacy, we need a few additional modifications for utility analysis.

Initialization. Random initialization has worked well for our empirical study. For the utility analysis, we need \hat{V}^0 to be reasonably close to V^* (in terms of spectral norm). We show that by using the noisy power iteration process, we are able to obtain \hat{V}^0 within the required bound with n almost linear in m , an improvement compared to private PCA (Dwork et al., 2014), which would require $n = \tilde{\Omega}(m\sqrt{m})$. See Section 4.1 for details.

Sampling from Ω . To ease the analysis, we require the observed values to be independent across different steps. This is achieved by resampling from Ω at the beginning of $\mathcal{A}_{\text{item}}$ (Line 6) and $\mathcal{A}_{\text{user}}$ (Line 13). The sampling in $\mathcal{A}_{\text{item}}$ is more important as it also limits the number of items per user, for privacy purposes. In practice, we omit the sampling in $\mathcal{A}_{\text{user}}$, and sample only once for $\mathcal{A}_{\text{item}}$. The sampling distribution used in the latter has a significant impact in practice, as discussed in Section 5.1.

3.2. Computational Complexity

The computational complexity of DPALS is comparable to that of ALS. More precisely, the V step of ALS involves computing \mathbf{H}_j^t and \mathbf{w}_j^t , in $O(|\Omega'|r^2)$, then solving the m linear systems $\hat{\mathbf{V}}_j^{t+1} = (\mathbf{H}_j^t)^{-1} \mathbf{w}_j^t$ in $O(mr^3)$, for a total complexity of $O(|\Omega'|r^2 + mr^3)$ (and similarly for the U step). This scales linearly in the number of observations $|\Omega'|$ and the number of items m . In the private version ($\mathcal{A}_{\text{user}}$), the only additional operations are forming the noise matrices (Lines 7–8) in $O(mr^2)$, and projecting \mathbf{X}_j (Line 10), in $O(mr^3)$, so the complexity per iteration is the same as ALS. In comparison, the complexity of DPFW is $O(\Gamma(m + |\Omega'|))$, using Oja’s method. The per-iteration complexity also scales linearly in m and $|\Omega'|$. Even though the per-iteration complexity of DPFW and DPALS are comparable, DPALS converges in much fewer iterations (see Appendix D.4 for an example), which makes it more scalable in practice.

3.3. Privacy Guarantee

We now provide the privacy guarantee for DPALS. As each subroutine in DPALS is a variant of the Gaussian mechanism, we can apply the Rényi accounting (Mironov, 2017) and convert to (ε, δ) -DP. See Appendix A for the proof.

Theorem 1 (Privacy guarantee). *Excluding the initialization of $\hat{\mathbf{V}}^0$, Algorithm 1 is $(\alpha, \alpha\rho^2)$ -joint RDP with $\rho^2 = \frac{kT}{2\sigma^2}$. Hence for any $\varepsilon > 0$ and $\delta \in (0, 1)$, Algorithm 1 is (ε, δ) -joint DP if we set $\sigma = \frac{\sqrt{(2kT)(\varepsilon + \log(1/\delta))}}{\varepsilon}$.*

The guarantee holds for all values of the parameters $\Gamma_u, \Gamma_M, T, \lambda, r, k$. Note in particular that the scale of the noise (Lines 7–8 in Algorithm 1) is normalized so that the expression of σ in Theorem 1 does not depend on Γ_u, Γ_M .

In the above guarantee, we have excluded the initialization process. With random initialization, which we use in practice, there is no extra privacy cost. However, for the utility guarantee, we need an extra DP procedure, detailed in Section 4.1, so that $\hat{\mathbf{V}}^0$ is sufficiently close to \mathbf{V}^* . This can be done within the same privacy bound as the main procedure.

4. Convergence Guarantee for DPALS

We now show that under standard low-rank matrix completion assumptions (Assumptions 1 and 2), Algorithm 1, initialized with the noisy power method, solves the matrix completion problem. We will first present the results assuming that $\hat{\mathbf{V}}^0$ is close to \mathbf{V}^* . We will then present the guarantee of the initialization procedure.

Theorem 2 (Utility guarantee). *Suppose that \mathbf{M} is a μ -incoherent rank- r matrix, and Ω consists of random observations with probability p . Let $\sigma_1^* \geq \dots \geq \sigma_r^* > 0$ be the*

singular values of \mathbf{M} and $\kappa := \sigma_1^/\sigma_r^*$ its condition number.*

There exists a universal constant $C > 0$, such that for all $\delta \in (0, 1)$, $\varepsilon \in (0, \log(1/\delta))$, if $p \geq \mu^6 \kappa^{12} r^6 \cdot \frac{\log^3 n}{m}$ and $\sqrt{pn} \geq C \frac{\gamma \sqrt{\log(1/\delta)}}{\varepsilon}$, where $\gamma = C \kappa^6 \mu^3 r^2 \sqrt{m} \cdot \log^2(\kappa \cdot n)$, then DPALS, initialized with $\hat{\mathbf{V}}^0$ s.t. $\|(I - \mathbf{V}^(\mathbf{V}^*)^\top) \hat{\mathbf{V}}^0\| \leq \frac{C}{\kappa^2 r^2 \log n}$, with parameters $k = C \cdot m \cdot p \log n$, $T = \log(\mu \kappa n / \varepsilon)$, $\sigma = \frac{C \sqrt{kT \log(1/\delta)}}{\varepsilon}$, $\Gamma_u = \frac{C \mu \sigma_1^* \sqrt{r}}{\sqrt{n}}$, $\Gamma_M = \frac{\mu^2 r \sigma_1^*}{\sqrt{mn}}$ and $\lambda = 0$, returns $\hat{\mathbf{U}}^T$ and $\hat{\mathbf{V}}^T$ such that the following holds:*

- *The distribution of $(\hat{\mathbf{U}}^T, \hat{\mathbf{V}}^T)$ satisfies (ε, δ) -joint DP.*
- *$\|\mathbf{M} - \hat{\mathbf{U}}^T (\hat{\mathbf{V}}^T)^\top\|_F \leq C \cdot \frac{\sqrt{m \log(1/\delta)}}{\varepsilon \cdot n} \cdot \frac{\kappa \gamma}{\sqrt{p}} \|\mathbf{M}\|_F$, with probability $\geq 1 - 1/n^{10}$.*
- *Similarly, $\|\mathbf{M} - \hat{\mathbf{U}}^T (\hat{\mathbf{V}}^T)^\top\|_\infty \leq C \cdot \frac{\sqrt{m \log(1/\delta)}}{\varepsilon \cdot n} \cdot \frac{\kappa \gamma}{\sqrt{p}} \cdot \frac{\mu^2 r \|\mathbf{M}\|_2}{\sqrt{mn}}$, with probability $\geq 1 - 1/n^{10}$.*

Remark 1. The choice of hyper-parameters in Theorem 2 assumes knowledge of certain quantities such as r, μ, κ . In practice, these quantities are unknown, but one can use standard DP hyper-parameter search techniques (Liu & Talwar, 2019) to search for optimal hyper-parameter values.

Remark 2. The number of samples needed per user is about $p \cdot m = O(\mu^6 \kappa^{12} r^6 \log^3 n)$ which is nearly optimal with respect to m and n . This represents a significant improvement over the DPFW algorithm in (Jain et al., 2018) which requires $\Omega(\sqrt{m})$ samples per user.

Remark 3. We did not optimize bounds for dependence on the rank r and condition number κ . Prior work tends to focus on the dependence on the size (m and n) and polynomial dependence on r, κ is common even in the non-private setting, for example (Jain et al., 2013; Sun & Luo, 2015; Ge et al., 2016). Our main goal is to provide a guarantee in the private setting that is competitive with the non-private setting, so we inherit the focus on the size m, n . Furthermore, dependence on κ can be removed (up to log factors) by using a *stagewise* ALS method similar to (Hardt & Wooters, 2014). However, this further complicates the proof and the practical performance of standard ALS is comparable to such stagewise methods.

Remark 4. Our Frobenius norm error bound is significantly smaller than the bound for the DPFW algorithm, which is given by $\|\mathbf{M} - \hat{\mathbf{U}}^T (\hat{\mathbf{V}}^T)^\top\|_F \leq \left(\frac{m^{5/4}}{n\varepsilon}\right)^{1/5} \|\mathbf{M}\|_F$. In particular, to ensure an error $\|\mathbf{M} - \hat{\mathbf{U}}^T (\hat{\mathbf{V}}^T)^\top\|_F \leq \zeta \|\mathbf{M}\|_F$, DPALS requires $n \geq \frac{Cm}{\zeta \cdot \varepsilon}$, while DPFW requires $n \geq \frac{Cm^{5/4}}{\zeta^{5 \cdot \varepsilon}}$, which is significantly worse in terms of ζ . Furthermore, the DPFW bound is a generalization bound, i.e., there is an additional bias term which can be large, and to the best of our knowledge, existing techniques (even in the non-private setting) require incoherence to control this term.

Remark 5. Consider a set of m linear regression problems in r -dimensions: $\{\mathbf{y}_{(i)} = \mathbf{X}\theta_{(i)}^*\}_{i=1}^m$, with $\mathbf{X} \in \mathbb{R}^{n \times r}$.

One can use a single iteration of DPALS with $(\widehat{\mathbf{U}} = \mathbf{X}$ and $\mathbf{P}_\Omega(\mathbf{M}) = [\mathbf{y}_{(1)}, \dots, \mathbf{y}_{(m)}])$ to solve these linear regression problems. Assuming the conditions on \mathbf{M} are satisfied, we can obtain an excess empirical risk of $\widehat{O}(\sqrt{m}/(\varepsilon n))$. This matches the best known upper bound for solving a set of linear regressions with privacy (Sheffett, 2019; Smith et al., 2017). So, a better convergence rate of DPALS would lead to a tighter bound on solving a set of linear regressions with a common feature matrix. For $m = O(1)$, we know that the lower bound for private linear regression is $\widehat{\Omega}(1/\varepsilon n)$ (Smith et al., 2017). Thus, we conjecture that the error for DPALS is tight w.r.t. m and εn .

Remark 6. Instead of using the perturbed objective function to estimate $\widehat{\mathbf{V}}^t$ in DPALS, one can use DPSPGD (Bassily et al., 2014) to do the same (solving a least squares problem with $\widehat{\mathbf{U}}^t$ fixed). We leave the empirical comparison of this approach to future work. However, we know that for least-square losses, perturbing the objective is known to be theoretically optimal (Smith et al., 2017).

Proof sketch: First, we show that under the assumptions in Theorem 2, w.h.p., clipping and sampling operations in DPALS have no effect. Note, using $k \geq Cp \cdot m \log n$, w.p. $\geq 1 - 1/n^{100}$, $\forall i, |\Omega_i| \leq k$. Furthermore, using Lemma 3, $\|\widehat{\mathbf{U}}_i^t\| \leq \Gamma_u$. Similarly, using Lemma 3, $\sigma_{\min}(\mathbf{X}) \geq p/4 - \|\mathbf{G}\|_2 \geq p/4 - \Gamma_u^2 \sigma \sqrt{r} \geq p/8$. That is, $\mathbf{X} \succ 0$.

The above observation implies that, under the assumptions of the theorem, Algorithm 1 is essentially performing the following iterative steps:

- i) $\widehat{\mathbf{U}}^t = \arg \min_{\widehat{\mathbf{U}}} \|\mathbf{P}_\Omega(\mathbf{M} - \widehat{\mathbf{U}}(\widehat{\mathbf{V}}^t)^\top)\|_F^2$, and
- ii) $\widehat{\mathbf{V}}_j^{t+1} = \left(\mathbf{I} + \sum_{i \in \Omega_j^t} \widehat{\mathbf{U}}_i^t \otimes \widehat{\mathbf{U}}_i^t + \mathbf{G} \right)^{-1} \left(\sum_{i \in \Omega_j^t} \mathbf{M}_{ij} \widehat{\mathbf{U}}_i^t + \mathbf{g} \right)$.

Let \mathbf{U}^t (resp. \mathbf{V}^t) be the Q part in the QR decomposition of $\widehat{\mathbf{U}}^t$ (resp. $\widehat{\mathbf{V}}^t$). Using Lemma 4, we get $\text{Err}(\mathbf{V}^*, \mathbf{V}^{t+1}) \leq \frac{1}{4} \text{Err}(\mathbf{V}^*, \mathbf{V}^t) + \alpha$, where $\text{Err}(\mathbf{V}^*, \mathbf{V}) = \|(\mathbf{I} - \mathbf{V}^*(\mathbf{V}^*)^\top) \mathbf{V}\|_F$ and $\alpha \leq \frac{C\kappa^6 \cdot \mu^3 r^2 \sqrt{\log n} \sqrt{m \log n} \cdot T \log 1/\delta}{\sqrt{pn} \varepsilon}$.

That is, after T iterations, $\text{Err}(\mathbf{V}^*, \mathbf{V}^T) \leq 2\alpha$. The second claim of the theorem now follows from the above observation and Lemma 3. Similarly, the third claim follows by using the bound on $\text{Err}(\mathbf{V}^*, \mathbf{V}^T)$ and incoherence of \mathbf{U}^T , \mathbf{V}^T (Lemma 3). See Appendix B for a detailed proof.

Lemma 3. Suppose the assumptions of Theorem 2 hold. Then, w.p. $\geq 1 - 5T/n^{100}$, we have: a) each iterate $\widehat{\mathbf{U}}^t$, $\widehat{\mathbf{V}}^t$ is $16\kappa\mu$ -incoherent, b) $1/2 \leq \sigma_q(\widehat{\mathbf{U}}^t(\boldsymbol{\Sigma}^t)^{-1}) \leq 2$ for all $q \in [r]$, c) $1/4 \leq \sigma_q(\frac{1}{p} \sum_{i:(i,j) \in \Omega^{v,t}} \widehat{\mathbf{u}}_i^t (\widehat{\mathbf{u}}_i^t)^\top) \leq 4$.

Lemma 4. Suppose the assumptions of Theorem 2 hold. Also, let \mathbf{V}^t be $16\kappa\mu$ -incoherent s.t. $\text{Err}(\mathbf{V}^*, \mathbf{V}^t) \leq$

$\frac{1}{\kappa^2 \log^2 n}$. Then, w.p. $\geq 1 - 5T/n^{100}$, we have $\text{Err}(\mathbf{U}^*, \mathbf{U}^t) \leq \frac{1}{2} \text{Err}(\mathbf{V}^*, \mathbf{V}^t)$, and $\text{Err}(\mathbf{V}^*, \mathbf{V}^{t+1}) \leq \frac{1}{2} \text{Err}(\mathbf{U}^*, \mathbf{U}^t) + \frac{C\kappa^6 \cdot \mu^3 r^2 \sqrt{\log n} \sqrt{m \log n} \cdot T \log 1/\delta}{\sqrt{pn} \varepsilon}$, where $\text{Err}(\mathbf{V}^*, \mathbf{V}) = \|(\mathbf{I} - \mathbf{V}^*(\mathbf{V}^*)^\top) \mathbf{V}\|_F$.

4.1. Noisy Power Iteration Initialization

Theorem 2 requires that DPALS be initialized with $\widehat{\mathbf{V}}^0$ such that $\|(\mathbf{I} - \mathbf{V}^*(\mathbf{V}^*)^\top) \widehat{\mathbf{V}}^0\| = O(1/\log n)$. One may apply Algorithm 1 of (Dwork et al., 2014), i.e. compute the top- r eigenvectors of $\mathbf{A} + \mathbf{G}$, where $\mathbf{A} := \mathbf{P}_\Omega(\mathbf{M})^\top \mathbf{P}_\Omega(\mathbf{M})$ and $\mathbf{G} \sim \mathcal{N}_{\text{sym}}(0, \Gamma_M^4 \sigma^2)^{m \times m}$. This would require $n = \widehat{\Omega}(m\sqrt{m}/\varepsilon)$. However, this turns out to be suboptimal in our setting as it doesn't take advantage of the sparsity of $\mathbf{P}_\Omega(\mathbf{M})$. Instead, the noisy power iteration method, developed in (Hardt & Roth, 2012; 2013; Hardt & Price, 2013) for per-entry privacy protection, turns out to be more suitable.

One difficulty in applying noisy power iteration is that prior work requires incoherence of \mathbf{A} , which may not hold in our setting. To overcome this difficulty, we show that it suffices to have incoherence of the top- r eigenspace of \mathbf{A} , together with a (moderate) gap between the top eigenvalues and the rest, both of which we are able to establish. This gives a tighter analysis of noisy power iteration which may be of independent interest, detailed in Appendix B.4. We apply this result to our setting in the next theorem.

Theorem 5 (Initialization guarantee). There exists constant $C_0, C_1, C_2 > 0$, such that for any $\delta \in (0, 1)$, $\varepsilon \in (0, \log(1/\delta))$, if $p \geq C_1 \gamma \log^3 m/m$ and $\sqrt{pn} \geq C_2 \frac{\sqrt{\log(1/\delta)}}{\varepsilon} \gamma \sqrt{m} \log^{5/2} m$, where $\gamma = (\mu\kappa r)^{C_0}$, the noisy power iteration method is (ε, δ) -differentially private, and with high probability returns a $\widehat{\mathbf{V}}^0$ which is close to \mathbf{V}^* as defined in Theorem 2.

5. Heuristic Improvements to DPALS

We introduce heuristics to improve the privacy/utility trade-off of DPALS in practice. We describe each heuristic, its motivation, and how to implement it differentially privately.

5.1. Reducing Distribution Skew

The first heuristics are motivated by the observation that, in practice, the elements of Ω are not sampled uniformly at random (Marlin et al., 2007). In particular, the number of observed ratings per item typically follows a power-law distribution, and is heavily skewed towards popular items. For example, Figure 2 shows the fraction of observations vs. fraction of top movies in the MovieLens 10M data set. It shows, for instance, that the top 20% of the movies account for more than 85% of the observations.




Figure 2. Fraction of observations contributed by the top movies in MovieLens 10M. Adaptive sampling reduces popularity bias.

Due to this popularity bias, some items may have very few observations, and for such rare items j , the embedding \mathbf{V}_j learned by DPALS may not be useful: The noise terms in Line 10 of Algorithm 1 do not scale with the number of observations $|\Omega'_j|$ – for otherwise we may lose the protection on users who rated rare items – thus, items with a smaller $|\Omega'_j|$ have a lower signal-to-noise ratio. In our experiments, we found that such noisy embeddings may have a further cascading effect and lead to quality degradation in the embeddings of other movies and users. To alleviate this issue, we propose two techniques.

Learning on frequent items. The first strategy is to partition the items into two sets, based on an estimate of the item counts, which we denote by $\tilde{c} \in \mathbb{R}^n$. We introduce a hyper-parameter β representing the fraction of movies to train on. Define the set Frequent to be the $\lceil m\beta \rceil$ items with the largest \tilde{c} , and let Infrequent be its complement. We learn embeddings $\hat{\mathbf{V}}_j$ only for $j \in \text{Frequent}$, by running Algorithm $\mathcal{A}_{\text{DPALS}}$ on those items. When making predictions for any missing entry M_{ij} , if $j \in \text{Frequent}$, we use the dot product $\hat{\mathbf{U}}_i \cdot \hat{\mathbf{V}}_j$, and if $j \in \text{Infrequent}$ we use the average observed rating of $P_\Omega(\mathbf{M})_i$.

To compute \tilde{c} privately, notice that since each user contributes at most k items, the exact item count c has ℓ_2 sensitivity \sqrt{k} . Thus, $\tilde{c} := c + \mathcal{N}(0, k\sigma^2)$ guarantees $(\alpha, \alpha/2\sigma^2)$ -RDP.

Adaptive sampling. To further reduce the popularity bias, we propose to use an adaptive distribution when subsampling Ω . Recall that in Line 6 of $\mathcal{A}_{\text{item}}$, we pick k items per user in Ω , in order to limit the privacy loss. We propose to sample rare items with higher probability, as follows. Given the count estimate \tilde{c} , for each user i , we pick the k items in $\Omega_i \cap \text{Frequent}$ with the lowest count estimates. This heuristic effectively reduces the distribution skew and gives a significant utility gain compared to uniform sampling, see Section 6.3. Figure 2 illustrates the resulting distribution for a sample size of $k = 50$ per user. It’s interesting to observe that under uniform sampling, the popularity bias is worse than in the unsampled data set, this is due to a negative correlation between user counts and item counts:

conditioned on a light user, the probability to observe a rare item is lower; see Appendix C for further discussion.

5.2. Additional Heuristics

A common heuristic, used for example by (McSherry & Mironov, 2009), is to center the observed matrix $P_\Omega(\mathbf{M})$, by subtracting an estimate of the global average, denoted by \tilde{m} . To compute \tilde{m} privately, since $\|P_\Omega(\mathbf{M})\|_\infty \leq \Gamma_M$ and each user contributes at most k items, publishing $\tilde{m} = \frac{\sum_{(i,j) \in \Omega} M_{ij} + \mathcal{N}(0, k\Gamma_M^2\sigma^2)}{|\Omega| + \mathcal{N}(0, k\sigma^2)}$ guarantees $(\alpha, \alpha/\sigma^2)$ -RDP.

Another practice, commonly used in some benchmarks, is to modify the loss function in Section 2.2 by adding the term $\lambda_0 \|\hat{\mathbf{U}}\hat{\mathbf{V}}^\top\|_F^2$, where λ_0 is a hyper-parameter. This is particularly important for item recommendation tasks, such as the MovieLens 20M benchmark. This modification introduces an additional term $\mathbf{K} := \lambda_0 \sum_{i \in [n]} \hat{\mathbf{U}}_i \otimes \hat{\mathbf{U}}_i$ to \mathbf{X} in Line 9 of $\mathcal{A}_{\text{item}}$. To maintain privacy, we use a noisy version $\tilde{\mathbf{K}}$ obtained by adding Gaussian noise to \mathbf{K} . Since \mathbf{K} is independent of j , we reuse the same $\tilde{\mathbf{K}}$ for all $j \in [m]$, thus limiting the additional privacy loss due to this term.

Finally, we project the matrix $\mathbf{X}_j = \mathbf{H}_j + \mathbf{G}_j$ to the PSD cone (Line 9) to improve stability. In our analysis, we show that \mathbf{X}_j is positive definite with high probability, but in practice, the projection improves performance.

We account for the privacy cost in the computation of \tilde{m} , \tilde{c} , and $\tilde{\mathbf{K}}$, along with that in Theorem 1, by standard composition properties of RDP (Mironov, 2017). For completeness, the privacy accounting of the full algorithm including data pre-processing, is given in Appendix C.

6. Empirical Evaluation

We run experiments on synthetic data and two benchmark tasks on the widely used MovieLens data sets (Harper & Konstan, 2016). The synthetic task follows the assumptions of our theoretical analysis, and serves to illustrate the guarantees of Theorem 2. The MovieLens benchmark tasks serve as an evaluation of the empirical privacy/utility trade-off on a more realistic application, and to provide some practical insights into DPALS. We use current SOTA method DPFW as the main baseline as it is already demonstrated to be more accurate than techniques like Private SVD (McSherry & Mironov, 2009). Similar to (Jain et al., 2018), we do not compare against (Liu & Talwar, 2019) as the privacy parameters are unclear, and might require (exponential time) Markov chain based sampling methods to compute them.

6.1. Metrics and Data Sets

Metrics. The quality of a learned model $(\hat{\mathbf{U}}, \hat{\mathbf{V}})$ will be measured either using the RMSE or the Recall@k, de-




Figure 3. Privacy/utility trade-off of different methods. We observe that DPALS is significantly more accurate than DPFW method, and the loss in accuracy for DPALS compared to ALS is relatively small, especially for $\epsilon \geq 10$.




Figure 4. Comparison of DPFW and DPALS on synthetic data with different number of rows/users n .

pending on the benchmark. The RMSE is defined as $\text{RMSE} = \|\mathbb{P}_{\Omega^{\text{test}}}(\hat{\mathbf{U}}\hat{\mathbf{V}}^{\text{T}} - \mathbf{M})\|_F / \sqrt{|\Omega^{\text{test}}|}$, where Ω^{test} is the set of test ratings held out from Ω . Recall@ k is defined as follows. For each user i , let R_i be the set of k movies with the highest scores, where the score of movie j is $\hat{\mathbf{U}}_i \cdot \hat{\mathbf{V}}_j$. Then $\text{Recall}@k = \frac{1}{n} \sum_{i=1}^n |R_i \cap \Omega_i^{\text{test}}| / \min(k, |\Omega_i^{\text{test}}|)$.

Synthetic data. We generate a rank 5 ground truth matrix as the product of two random orthogonal matrices $\mathbf{U}^* \in \mathbb{R}^{n \times 5}$, $\mathbf{V}^* \in \mathbb{R}^{m \times 5}$, where $m = 1000$, and $n \in \{5000, 10000, 20000, 50000\}$. We scaled the ground truth matrix such that the standard deviation of the observations is 1, in other words, a trivial model which always predicts the global average has a RMSE of 1. The observed entries Ω are obtained by sampling each entry independently with probability $p = 20 \log(n)/m$.

MovieLens data sets. We apply our method to two common recommender benchmarks: (i) rating prediction on MovieLens 10M (ML-10M) following Lee et al. (2013), where the task is to predict the value of a user’s rating, and performance is measured using the RMSE, (ii) item recommendation on MovieLens 20M (ML-20M) following Liang et al. (2018), where the task is to select k movies for each user and performance is measured using Recall@ k . For comparison to DPFW, we use a variant of the ML-10M task following Jain et al. (2018), where the movies are restricted

to the 400 most popular movies (DPFW did not scale to the full data set with all movies, unlike DPALS).

Experimental protocol. Each data set is partitioned into training, validation and test sets. Hyper-parameters are chosen on the validation set, and the final performance is measured on the test set. The privacy loss accounting is done using RDP, then translated to (ϵ, δ) -DP with $\delta = 10^{-5}$ for the synthetic data and ML-10M and $\delta = 1/n$ for ML-20M. When training DPALS models on synthetic data, we use the basic Algorithm 1, without heuristics. When training on MovieLens, we use the heuristics described in Section 5. Note that even when training on Frequent items (Section 5.1), evaluation is always done on the *full set* of items, so that the reported metrics are comparable to previously published numbers. Additional details on the experimental setup are in Appendix D, including a list of hyper-parameters and the ranges we used for each.

6.2. Privacy-Utility Trade-Off

DPALS vs. DPFW on synthetic data. On synthetic data (Figure 4) we observe: First, as expected, the trade-off of both algorithms improves as the number of users increases. Second, for $\epsilon = 1$, the quality of the DPFW models is no better than the trivial model (RMSE equal to 1), while DPALS has a lower RMSE, which significantly improves with larger n . Third, for the largest data set ($n = 50K$), the relative improvement in RMSE between DPALS and DPFW is at least 7-fold across all values of ϵ . To further illustrate the difference between DPALS and DPFW, we show in Appendix D.4 the RMSE against number of iterations, both for the private and non-private variants (Figure 7).

DPALS vs. DPFW on ML10M. Next, we compare the two methods on ML-10M-top400 (Figure 3a). For DPFW and DPSVD, the numbers are taken directly from (Jain et al., 2018). For reference, we include the test RMSE of non-private ALS, and a simple baseline model that always pre-

dicts the global average rating. The performance of DPSVD is worse than that of the simple baseline. DPALS performs best, with a relative improvement in RMSE (compared to DPFW) that ranges from 7% to 11.6%, and that increases with ϵ . In Appendix D.4, we show that DPALS achieves performance better than DPFW even when trained on a small fraction of the users (23%).

Finally, Figure 3b shows the privacy/utility trade-off on the full ML-10M data. In order to scale DPFW to the full data, we use the same procedure described in Section 5: DPFW is trained on the top movies, and for remaining movies the model predicts the user’s average rating. Compared to the restricted data set (ML-10M-top400), the privacy-utility trade-off is worse on the full data. This indicates that a smaller ratio between number of users and number of items makes the task harder – a result that is in line with the theory.

The results on synthetic data and ML10M suggest that DPALS exhibits a much better privacy/utility trade-off than DPFW, and a better dependence on the number of rows n , which is consistent with the theoretical analysis.

DPALS on MovieLens 20M. Figure 3c shows the privacy/utility trade-off of DPALS on the ML-20M data set. We include as a reference the non-private ALS, and a simple baseline model that always returns the k most rated movies.

On this task, the performance of the private model is remarkably good. Indeed, the best previously reported Recall@20 numbers for *non-private* models on this benchmark are 36.0% for ALS (Liang et al., 2018) and 41.4% using a sophisticated auto-encoder model (Shenbin et al., 2020). Our results show that DPALS can achieve performance comparable to the previously reported state of the art numbers for (non-private) matrix completion, and the utility does not significantly degrade, even at small ϵ .

6.3. Importance of Adaptive Sampling and Projection

In this section, we give additional insights into the effect of the heuristics introduced in Section 5. We run a study on ML-10M for $\epsilon = 10$, $r = 128$ and a sample size $k = 50$ (both correspond to the best overall model); other hyper-parameters are re-tuned. According to Section 5.1, we partition the set of movies into Frequent and Infrequent and train only on Frequent. The results are reported in Figure 5, where the movie fraction is simply $|\text{Frequent}|/n$. We make the following observations. First, for non-private ALS, we get the highest RMSE by training on all movies, while there is a benefit for training on a subset of the movies for the private models. Second, when training the non-private model on sub-sampled data (red and purple lines), there is a considerable increase in RMSE, from 0.785 to 0.812. This gives an indication that part of the utility loss is due to sub-

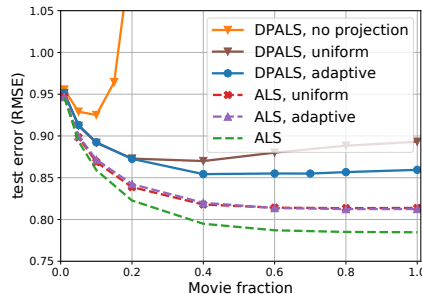


Figure 5. RMSE vs. movie fraction for $\epsilon = 10$ on ML-10M.

sampling, and not simply due to the addition of noise. Third, the sampling strategy has a significant impact on the performance of the private DPALS model: adaptive sampling improves the RMSE from 0.870 to 0.854, in contrast, the sampling strategy appears to have little effect on non-private models (i.e. models trained without noise). Finally, training the private model without PSD projection (Π_{PSD} in Line 10 of Algorithm 1) results in a terrible performance. We find that while the projection is not technically necessary for the theoretical analysis, it is essential in practice.

Training on a subset of the movies appears to have only a marginal effect when combined with adaptive sampling. However, as detailed in the appendix, the effect is much more significant for smaller ϵ , as well as on ML-20M.

Additional experiments are presented in Appendix D, to explore the effect of other hyper-parameters, such as the rank and the regularization of the objective function.

7. Conclusion

We presented DPALS for solving low-rank matrix completion with user-level privacy protection. We show that DPALS provably converges to high accuracy outputs under standard assumptions and, with careful implementation, significantly outperforms existing privacy preserving matrix completion methods. In fact, DPALS achieves competitive metrics on benchmark data compared to non-private models and scales well with data set size.

The efficiency of DPALS shows that by taking advantage of the structure of the problem, one can achieve a much higher utility for privacy-preserving model training. In this case, the alternating structure of ALS, along with the decoupling of the least squares solution, were essential in the design of an efficient method. These insights may be applicable to a broader class of problems and optimization algorithms.

Acknowledgments

We would like to thank Om Thakkar and the anonymous reviewers for insightful comments and discussion.

References

- Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.
- Bassily, R., Smith, A., and Thakurta, A. Private empirical risk minimization: Efficient algorithms and tight error bounds. In *Proc. of the 2014 IEEE 55th Annual Symp. on Foundations of Computer Science (FOCS)*, 2014.
- Bhatia, R. *Matrix analysis*, volume 169. Springer Science & Business Media, 2013.
- Calandrino, J. A., Kilzer, A., Narayanan, A., Felten, E. W., and Shmatikov, V. “you might also like:” privacy risks of collaborative filtering. In *2011 IEEE symposium on security and privacy*, pp. 231–246. IEEE, 2011.
- Candès, E. J. and Recht, B. Exact matrix completion via convex optimization. *Foundations of Computational mathematics*, 9(6):717–772, 2009.
- Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., and Song, D. The secret sharer: Evaluating and testing unintended memorization in neural networks. In *28th USENIX Security Symposium (USENIX Security 19)*, pp. 267–284, 2019.
- Carlini, N., Deng, S., Garg, S., Jha, S., Mahloujifar, S., Mahmood, M., Song, S., Thakurta, A., and Tramer, F. An attack on instahide: Is private learning possible with instance encoding? *arXiv preprint arXiv:2011.05315*, 2020a.
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., et al. Extracting training data from large language models. *arXiv preprint arXiv:2012.07805*, 2020b.
- Dekel, Y., Lee, J. R., and Linial, N. Eigenvectors of random graphs: Nodal domains. *Random Structures & Algorithms*, 39(1):39–58, 2011.
- Dinur, I. and Nissim, K. Revealing information while preserving privacy. In *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 202–210, 2003.
- Dwork, C. and Roth, A. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- Dwork, C., Kenthapadi, K., McSherry, F., Mironov, I., and Naor, M. Our data, ourselves: Privacy via distributed noise generation. In *Advances in Cryptology—EUROCRYPT*, pp. 486–503, 2006a.
- Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Proc. of the Third Conf. on Theory of Cryptography (TCC)*, pp. 265–284, 2006b.
- Dwork, C., McSherry, F., and Talwar, K. The price of privacy and the limits of lp decoding. In *Proceedings of the thirty-ninth annual ACM Symposium on Theory of Computing*, pp. 85–94, 2007.
- Dwork, C., Talwar, K., Thakurta, A., and Zhang, L. Analyze gauss: optimal bounds for privacy-preserving principal component analysis. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pp. 11–20, 2014.
- Erdős, L., Knowles, A., Yau, H.-T., and Yin, J. Spectral statistics of erdős–rényi graphs i: Local semicircle law. *The Annals of Probability*, 41(3B):2279–2375, 2013.
- Ge, R., Lee, J. D., and Ma, T. Matrix completion has no spurious local minimum. *Advances in Neural Information Processing Systems*, pp. 2981–2989, 2016.
- Hardt, M. and Price, E. The noisy power method: A meta algorithm with applications. *arXiv preprint arXiv:1311.2495*, 2013.
- Hardt, M. and Roth, A. Beating randomized response on incoherent matrices. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pp. 1255–1268, 2012.
- Hardt, M. and Roth, A. Beyond worst-case analysis in private singular vector computation. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 331–340, 2013.
- Hardt, M. and Wootters, M. Fast matrix completion without the condition number. In *Conference on learning theory*, pp. 638–678. PMLR, 2014.
- Hardt, M., Meka, R., Raghavendra, P., and Weitz, B. Computational limits for matrix completion. In *Conference on Learning Theory*, pp. 703–725. PMLR, 2014.
- Harper, F. M. and Konstan, J. A. The movielens datasets: History and context. *Acm Transactions on Interactive Intelligent Systems (TiiS)*, 5(4):19, 2016.
- Hu, Y., Koren, Y., and Volinsky, C. Collaborative filtering for implicit feedback datasets. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining, ICDM '08*, pp. 263–272, 2008.
- Jain, P. and Netrapalli, P. Fast exact matrix completion with finite samples. In *Conference on Learning Theory*, pp. 1007–1034. PMLR, 2015.

- Jain, P., Netrapalli, P., and Sanghavi, S. Low-rank matrix completion using alternating minimization. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pp. 665–674, 2013.
- Jain, P., Thakkar, O. D., and Thakurta, A. Differentially private matrix completion revisited. In *International Conference on Machine Learning*, pp. 2215–2224. PMLR, 2018.
- Kearns, M., Pai, M., Roth, A., and Ullman, J. Mechanism design in large games: Incentives and privacy. In *Proceedings of the 5th conference on Innovations in theoretical computer science*, pp. 403–410, 2014.
- Koren, Y. and Bell, R. Advances in collaborative filtering. *Recommender systems handbook*, pp. 77–118, 2015.
- Koren, Y., Bell, R., and Volinsky, C. Matrix factorization techniques for recommender systems. *Computer*, 42(8): 30–37, 2009.
- Korolova, A. Privacy violations using microtargeted ads: A case study. In *2010 IEEE International Conference on Data Mining Workshops*, pp. 474–482. IEEE, 2010.
- Lee, J., Kim, S., Lebanon, G., and Singer, Y. Local low-rank matrix approximation. In *Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, ICML'13*, pp. II–82–II–90. JMLR.org, 2013.
- Liang, D., Krishnan, R. G., Hoffman, M. D., and Jebara, T. Variational autoencoders for collaborative filtering. *WWW '18*, pp. 689–698, 2018.
- Liu, J. and Talwar, K. Private selection from private candidates. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing*, pp. 298–309, 2019.
- Liu, Z., Wang, Y.-X., and Smola, A. Fast differentially private matrix factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pp. 171–178, 2015.
- Lu, S., Hong, M., and Wang, Z. PA-GD: On the convergence of perturbed alternating gradient descent to second-order stationary points for structured nonconvex optimization. In *Proceedings of the 36th International Conference on Machine Learning*, pp. 4134–4143, 2019.
- Marlin, B. M., Zemel, R. S., Roweis, S., and Slaney, M. Collaborative filtering and the missing at random assumption. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence, UAI'07*, pp. 267–275, Arlington, Virginia, USA, 2007. AUAI Press.
- McSherry, F. and Mironov, I. Differentially private recommender systems: Building privacy into the netflix prize contenders. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 627–636, 2009.
- Meng, X., Wang, S., Shu, K., Li, J., Chen, B., Liu, H., and Zhang, Y. Personalized privacy-preserving social recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Mironov, I. Rényi differential privacy. In *2017 IEEE 30th Computer Security Foundations Symposium (CSF)*, pp. 263–275. IEEE, 2017.
- Recht, B. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12(12), 2011.
- Rendle, S., Zhang, L., and Koren, Y. On the difficulty of evaluating baselines: A study on recommender systems. *CoRR*, abs/1905.01395, 2019.
- Rudelson, M. and Vershynin, R. Delocalization of eigenvectors of random matrices with independent entries. *Duke Mathematical Journal*, 164(13):2507–2538, 2015.
- Sheffet, O. Old techniques in differentially private linear regression. In *Algorithmic Learning Theory*, pp. 789–827. PMLR, 2019.
- Shenbin, I., Alekseev, A., Tutubalina, E., Malykh, V., and Nikolenko, S. I. Recvae: A new variational autoencoder for top-n recommendations with implicit feedback. In *Proceedings of the 13th International Conference on Web Search and Data Mining, WSDM '20*, pp. 528–536, 2020.
- Shokri, R., Stronati, M., Song, C., and Shmatikov, V. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 3–18. IEEE, 2017.
- Smith, A., Thakurta, A., and Upadhyay, J. Is interaction necessary for distributed private learning? In *2017 IEEE Symposium on Security and Privacy (SP)*, pp. 58–77. IEEE, 2017.
- Song, S., Chaudhuri, K., and Sarwate, A. D. Stochastic gradient descent with differentially private updates. In *2013 IEEE Global Conference on Signal and Information Processing*, pp. 245–248. IEEE, 2013.
- Sun, R. and Luo, Z. Guaranteed matrix completion via nonconvex factorization. In *FOCS*, 2015.
- Thakkar, O., Ramaswamy, S., Mathews, R., and Beaufays, F. Understanding unintended memorization in federated learning. *arXiv preprint arXiv:2006.07490*, 2020.

- Tropp, J. A. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.
- Vershynin, R. Introduction to the non-asymptotic analysis of random matrices. *arXiv preprint arXiv:1011.3027*, 2010.
- Vu, V. and Wang, K. Random weighted projections, random quadratic forms and random eigenvectors. *Random Structures & Algorithms*, 47(4):792–821, 2015.
- Zhu, Y. and Wang, Y.-X. Improving sparse vector technique with renyi differential privacy. *Advances in Neural Information Processing Systems*, 33, 2020.

A. Proof of Theorem 1

Proof. To prove the guarantee for Algorithm 1, it suffices to show the following claim: that at each time step $t \in [T]$, the computations of \mathbf{X} and $\sum_{i \in \Omega'_j} \mathbf{M}_{ij} \cdot \mathbf{U}_i$, for all $j \in [m]$ satisfy $(\alpha, \alpha \frac{k}{2\sigma^2})$ -RDP. One can then compose the privacy losses via simple Rényi composition (Mironov, 2017) to obtain the overall RDP-cost to be $(\alpha, \alpha \frac{kT}{2\sigma^2})$.

To prove the claim, notice that at each time step $t \in [T]$, there are m computations of \mathbf{X} and $\sum_{i \in \Omega'_j} \mathbf{M}_{ij} \cdot \mathbf{U}_i$. Since each user $i \in [n]$ can affect only k of those computations, by Gaussian mechanism (Dwork et al., 2006a; Mironov, 2017) and the generalization of standard composition property of RDP (Mironov, 2017, Proposition 1) to the joint RDP, we have the required guarantee.

We now translate joint-RDP to join-DP. By the first part of the theorem, Algorithm 1 is $(\alpha, \alpha\rho^2)$ -joint RDP with $\rho^2 = \frac{kT}{2\sigma^2}$. Thus by (Mironov, 2017, Proposition 3) it is (ε, δ) -joint DP with $\varepsilon = \alpha\rho^2 + \frac{\log(1/\delta)}{\alpha-1}$ for any $\alpha > 1$. The latter expression is minimized when $\alpha = 1 + \frac{\sqrt{\log(1/\delta)}}{\rho}$, which yields $\varepsilon^{\min}(\rho) = \alpha\rho^2 + \frac{\log(1/\delta)}{\alpha-1} = 2\sqrt{\log(1/\delta)}\rho + \rho^2$. Now fix $\varepsilon > 0, \delta \in (0, 1)$. To guarantee (ε, δ) -joint DP while minimizing the noise (which scales as $1/\rho$ by definition of ρ), it suffices to maximize ρ subject to $\varepsilon^{\min}(\rho) \leq \varepsilon$, but since ε^{\min} is increasing in ρ , ρ is maximized when $\varepsilon^{\min}(\rho) = \varepsilon$. This is a second-order polynomial in ρ , and it has a positive root at $\rho^+ = \sqrt{\log(1/\delta) + \varepsilon} - \sqrt{\log(1/\delta)}$. Therefore, setting

$$\sigma = \frac{\sqrt{kT/2}}{\rho^+} = \frac{\sqrt{kT/2}}{\sqrt{\log(1/\delta) + \varepsilon} - \sqrt{\log(1/\delta)}} = \frac{\sqrt{kT/2}(\sqrt{\log(1/\delta) + \varepsilon} + \sqrt{\log(1/\delta)})}{\varepsilon} \leq \frac{\sqrt{2kT(\log(1/\delta) + \varepsilon)}}{\varepsilon}$$

suffices to guarantee (ε, δ) -joint DP. This completes the proof. \square

B. Proofs from Section 4

Recall the problem setting. $\mathbf{M} = \mathbf{U}^* \mathbf{\Sigma}^* (\mathbf{V}^*)^\top$ where $(\mathbf{U}^*)^\top \mathbf{U}^* = \mathbf{I}$ and $(\mathbf{V}^*)^\top \mathbf{V}^* = \mathbf{I}$. Also, \mathbf{U}^* and \mathbf{V}^* are μ -incoherent by assumption. That is, $\|\mathbf{U}_i^*\|_2 \leq \mu\sqrt{r}/\sqrt{n}$ and $\|\mathbf{V}_j^*\|_2 \leq \mu\sqrt{r}/\sqrt{m}$. The set of observations is $\Omega = \{(i, j) \text{ s.t. } \delta_{ij} = 1\}$, where δ_{ij} are i.i.d. Bernoulli random variables with $\Pr[\delta_{ij} = 1] = p$. We sample a new set of observations before every update.

We now present a basic lemma.

Lemma 6. *Let \mathbf{U}^* and \mathbf{U}^t be μ and μ_1 -incoherent, orthonormal matrices where $\mu_1 \geq \mu$ and $n \cdot p \geq \mu\mu_1 r^2$. Then, the following holds for all $j \in [m]$ (w.p. $\geq 1 - m\beta$):*

$$\left\| \frac{1}{p} \sum_{i=1}^n \delta_{ij} \mathbf{U}_i^* (\mathbf{U}_i^t)^\top - (\mathbf{U}^*)^\top \mathbf{U}^t \right\|_F \leq C \sqrt{\frac{\mu_1^2 r}{n \cdot p}} \cdot \log \frac{r}{\beta}.$$

Proof. The proof follows from the matrix Bernstein inequality (Tropp, 2015, Theorem 6.1.1) and incoherence of $\mathbf{U}^*, \mathbf{U}^t$. \square

Lemma 7. *Let δ_{ij} be i.i.d. Bernoulli random variables with $\Pr[\delta_{ij} = 1] = p$. Then, the following holds (w.p. $\geq 1 - \delta$):*

$$\left\| \frac{1}{p} \mathbf{P}_\Omega(\mathbf{M}) - \mathbf{M} \right\|_F \leq C \left(\sqrt{\frac{n}{p} \log \frac{1}{\beta}} + \frac{1}{p} \log \frac{1}{\beta} \right) \cdot \|\mathbf{M}\|_\infty.$$

Proof. The lemma is similar to Theorem 7 of (Recht, 2011) and follows by the matrix Bernstein inequality (Tropp, 2015, Theorem 6.1.1). \square

B.1. Rank-1 Case

Simplifying the notation, denote $\mathbf{M} = \sigma^* \mathbf{u}^* (\mathbf{v}^*)^\top$ where $(\mathbf{u}^*)^\top \mathbf{u}^* = 1$ and $(\mathbf{v}^*)^\top \mathbf{v}^* = 1$.

Note that as $k = C \cdot p \cdot m \log n$ w.p. $\geq 1 - T/n^{100}$, we do not throw any tuples in Line 6 of Algorithm 1. Similarly, using incoherence we have: $\|\mathbf{M}\|_\infty \leq \Gamma_{\mathbf{M}}$. So, we do not clip any sample in Line 1 of Algorithm 1.

Now, we use mathematical induction to show the incoherence of resulting \mathbf{v}^t and $\hat{\mathbf{u}}^t$, and to show that the clipping operations do not really apply in our setting with the selected hyper-parameters.

For the base case ($t = 0$), initialization of \mathbf{v}^0 ensures that $\text{Err}(\mathbf{v}^*, \mathbf{v}^0) \leq \frac{C}{\log n}$. Now, using (Jain et al., 2013, Lemma C.2) that uses clipping only in the first step to ensure incoherence, we get that \mathbf{v}^0 is 16μ -incoherent.

In the induction step, assuming the Lemma holds for \mathbf{v}^t , we prove the claim for $\hat{\mathbf{u}}^t$ and \mathbf{v}^{t+1} . Dropping superscripts of Ω'_i for notation simplicity and using $\lambda = 0$, we have: $\hat{\mathbf{u}}^t = \arg \min_{\mathbf{u}} \|\text{P}_{\Omega}(M - \mathbf{u}(\mathbf{v}^t)^\top)\|_F^2$. The update of $\mathbf{u}^t = \hat{\mathbf{u}}^t / \|\hat{\mathbf{u}}^t\|_2$. So using (Jain et al., 2013, Lemma 5.5, Lemma 5.7, Theorem 5.1), we get w.p. $\geq 1 - 1/n^{100}$:

$$\begin{aligned} \mathbf{u}^t &\text{ is } 16\mu\text{-incoherent,} \\ \|\hat{\mathbf{u}}^t\|_2 &\geq \sigma^*/16, \\ \text{Err}(\mathbf{u}^t, \mathbf{u}^*) &\leq \frac{1}{4}\text{Err}(\mathbf{v}^t, \mathbf{v}^*). \end{aligned} \quad (3)$$

To complete the claim, we only need to study the update for \mathbf{v}^{t+1} , which is a noisy version of the ALS update:

$$\hat{\mathbf{v}}^{t+1} = (\mathbf{D} + \mathbf{G}^t)^{-1} (\text{P}_{\Omega}(M)\hat{\mathbf{u}}^t + \mathbf{g}),$$

where \mathbf{D} and \mathbf{G}^t are diagonal matrices s.t. $\mathbf{D}_{jj} = \sum_{(i,j) \in \Omega} (\hat{\mathbf{u}}_i^t)^2$ and $\mathbf{G}_{jj}^t \sim \Gamma_{\mathbf{u}}^2 \sigma \cdot \mathcal{N}(0, 1)$.

We first prove that $\mathbf{D} + \mathbf{G}^t$ is indeed invertible, and has lower-bounded smallest eigenvalue. Using Lemma 6, and $pn \geq \mu^2 \log n \log(1/\delta)$, we have w.p. $\geq 1 - m\delta$,

$$\frac{1}{p} \mathbf{D}_{jj} \geq \|\hat{\mathbf{u}}^t\|_2^2 \left(1 - \sqrt{\frac{1}{\log n}}\right).$$

Also, using maximum of Gaussians, we have w.p. $\geq 1 - m\beta$,

$$\frac{1}{p} \|\mathbf{G}^t\|_2 \leq \frac{\Gamma_{\mathbf{u}}^2 \sigma \sqrt{\log(n/\beta)}}{p} \leq \frac{\mu^2 (\sigma^*)^2 \sigma \sqrt{\log(n/\beta)}}{np} \leq \frac{\|\hat{\mathbf{u}}^t\|_2^2}{16 \times 256},$$

where the final inequality follows by the assumption on p .

So,

$$\|(\mathbf{D} + \mathbf{G}^t)^{-1}\|_2 \leq \frac{2}{p \cdot \|\hat{\mathbf{u}}^t\|_2^2}. \quad (4)$$

We now conduct error analysis for $\hat{\mathbf{v}}^{t+1}$:

$$\hat{\mathbf{v}}^{t+1} = \alpha \cdot \mathbf{v}^* - \mathbf{E},$$

where $\alpha = \frac{\sigma^* \cdot (\mathbf{u}^*)^\top \mathbf{u}^t}{\|\hat{\mathbf{u}}^t\|_2}$. Furthermore, for a matrix \mathbf{C} with $\mathbf{C}_{jj} = \sum_{(i,j) \in \Omega} \hat{\mathbf{u}}_i^t \mathbf{u}_i^*$, we have $\mathbf{E} = \mathbf{E}^1 + \mathbf{E}^2$ with

$$\mathbf{E}_j^1 = (\mathbf{D}_{jj} + \mathbf{G}_{jj}^t)^{-1} (\alpha \mathbf{D}_{jj} - \sigma^* \mathbf{C}_{jj}) \mathbf{v}_j^*, \quad \mathbf{E}_j^2 = (\mathbf{D}_{jj} + \mathbf{G}_{jj}^t)^{-1} (\alpha \mathbf{G}_{jj}^t \mathbf{v}_j^* - \mathbf{g}_j).$$

This step follows from the observation that $(\text{P}_{\Omega}(M)\hat{\mathbf{u}}^t)_j = \sigma^* \mathbf{C}_{jj} \mathbf{v}_j^*$. (We note that \mathbf{E} is a vector but we use upper case to be consistent with Section B.2.)

Note that $\mathbf{E}[\alpha \mathbf{D}_{jj} - \sigma^* \mathbf{C}_{jj}] = 0$. Furthermore, using incoherence of \mathbf{v}^* , $\|\hat{\mathbf{u}}^t\|_2 \geq \sigma^*/16$, and the Bernstein's inequality, we have:

$$\|\mathbf{E}^1\|_2 \leq \frac{1}{64} \text{Err}(\mathbf{u}^t, \mathbf{u}^*). \quad (5)$$

Now,

$$\|\mathbf{E}^2\|_2 \leq \frac{2\sqrt{\log n}}{p \cdot \|\hat{\mathbf{u}}^t\|_2^2} \cdot (16\Gamma_{\mathbf{u}}^2 \sigma + \Gamma_M \Gamma_{\mathbf{u}} \sigma \sqrt{m}) \leq \frac{C\mu^4 \sqrt{\log n}}{np} \cdot \sigma. \quad (6)$$

Using (5) and (6), $\|\widehat{\mathbf{v}}^{t+1}\|_2 \geq 3/4$.

Thus, we get:

$$\text{Err}(\mathbf{v}^*, \mathbf{v}^{t+1}) \leq \frac{1}{32} \text{Err}(\mathbf{u}^*, \mathbf{u}^{t+1}) + \frac{C\mu^4 \sqrt{\log n}}{np} \cdot \sigma.$$

Similarly, by incoherence of \mathbf{v}^* and using bound on \mathbf{E}_j^1 and \mathbf{E}_j^2 , we get:

$$\|\widehat{\mathbf{v}}^{t+1}\|_\infty \leq 3\mu.$$

Therefore

\mathbf{v}^{t+1} is 16μ -incoherent,

$$\text{Err}(\mathbf{v}^{t+1}, \mathbf{v}^*) \leq \frac{1}{32} \text{Err}(\mathbf{u}^t, \mathbf{u}^*) + \frac{C\mu^4 \sqrt{\log n}}{np} \cdot \sigma. \quad (7)$$

So, the inductive hypothesis holds. Furthermore, we get Theorem 2, by combining the error terms of \mathbf{u}^t and \mathbf{v}^{t+1} .

B.2. Rank- r Case

B.2.1. PROOF OF LEMMA 3

Note that as $k = C \cdot p \cdot m \log n$, w.p. $\geq 1 - T/n^{100}$, we do not throw any tuples in Line 6 of Algorithm 1. Similarly, using incoherence we have: $\|\mathbf{M}\|_\infty \leq \Gamma_M$. So, we do not clip any samples in Line 1 of Algorithm 1.

Now, we use mathematical induction to show the incoherence of resulting $\widehat{\mathbf{V}}^t$ and $\widehat{\mathbf{U}}^t$, and to show that the clipping operations do not really apply in our setting with the selected hyperparameters.

For the base case ($t = 0$), initialization of $\widehat{\mathbf{V}}^0$ ensures that $(\widehat{\mathbf{V}}^0)^\top \widehat{\mathbf{V}}^0 = \mathbf{I}$ and $\text{Err}(\mathbf{V}^*, \widehat{\mathbf{V}}^0) \leq \frac{C}{\kappa^2 r^2 \log n}$. Now, using (Jain et al., 2013, Lemma C.2) that uses clipping only in the first step to ensure incoherence, we get that $\widehat{\mathbf{V}}^0$ is $16\mu\sqrt{r}$ -incoherent.

For the induction step, assuming the Lemma holds for $\widehat{\mathbf{V}}^t$, we prove the claim for $\widehat{\mathbf{U}}^t$ and $\widehat{\mathbf{V}}^{t+1}$. Dropping superscripts of Ω'_i for notation simplicity and using $\lambda = 0$, we have: $\widehat{\mathbf{U}}^t = \arg \min_{\mathbf{U}} \|\text{P}_\Omega(\mathbf{M} - \mathbf{U}(\mathbf{V}^t)^\top)\|_F^2$. That is, the update of $\widehat{\mathbf{U}}^t = \mathbf{U}^t \mathbf{R}_U$, with \mathbf{U}^t being the Q part of QR-decomposition, is identical to the standard non-noisy ALS. So using (Jain et al., 2013, Lemma 5.5, Lemma 5.7, Theorem 5.1)¹, we get w.p. $\geq 1 - 1/n^{100}$,

$$\begin{aligned} \mathbf{U}^t &\text{ is } 16\kappa\mu\text{-incoherent,} \\ \|\Sigma^* \mathbf{R}_U^{-1}\|_2 &\leq 16\kappa, \text{ i.e., } \|\mathbf{R}_U^{-1}\|_2 \leq \|(\Sigma^*)^{-1}\|_2 \|\Sigma^* \mathbf{R}_U^{-1}\|_2 \leq 16\|(\Sigma^*)^{-1}\|_2 \kappa, \\ \text{Err}(\mathbf{U}^t, \mathbf{U}^*) &\leq \frac{1}{4} \text{Err}(\mathbf{V}^t, \mathbf{V}^*). \end{aligned} \quad (8)$$

That is, now to complete the claim we only need to study the update for $\widehat{\mathbf{V}}^{t+1}$, which is a noisy version of the ALS update.

Now consider,

$$\mathbf{X}_j^t = (\widehat{\mathbf{U}}^t)^\top \left(\sum_{i:(i,j) \in \Omega} \mathbf{e}_i \mathbf{e}_i^\top \right) \widehat{\mathbf{U}}^t + \mathbf{G}_j^t = p \cdot \mathbf{R}_U \left((\mathbf{U}^t)^\top \left(\frac{1}{p} \sum_{i:(i,j) \in \Omega} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{U}^t + \mathbf{N}_j^t \right) \mathbf{R}_U, \quad (9)$$

where \mathbf{G}^t is the noise added in Line 9 of Algorithm 1 at time step t , $\mathbf{D}_j^t = (\mathbf{U}^t)^\top \left(\frac{1}{p} \sum_{i:(i,j) \in \Omega} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{U}^t$ and $\mathbf{N}_j^t = \frac{1}{p} \mathbf{R}_U^{-1} \mathbf{G}_j^t \mathbf{R}_U^{-1}$. Note that using Gaussian eigenvalue bound (Vershynin, 2010) and Weyl's inequality (Bhatia, 2013), we have w.p. $\geq 1 - 1/n^{100}$,

$$\sigma_{\min}(\mathbf{D}_j^t + \mathbf{N}_j^t) \geq \left(1 - C \sqrt{\frac{\mu^2 \kappa^2 r}{n \cdot p}} \cdot \log n - \frac{2\Gamma_{\mathbf{u}}^2 \sigma \sqrt{r}}{p \cdot \sigma_{\min}(\mathbf{R}_U)^2} \right) \geq \frac{1}{2}, \quad (10)$$

¹Lemma 5.5 of (Jain et al., 2013) has a redundant \sqrt{r} term in incoherence claim

where the last inequality follows from: $np \geq C\mu^2\kappa^2r \log^2 n$ and $n\sqrt{p} \geq C\mu^2\kappa^6r\sqrt{r} \cdot \frac{\sqrt{m \log n} \cdot T \log(1/\delta)}{\varepsilon}$.

Next, we argue that \mathbf{X}_j^t is PSD. Observe that

$$\sigma_{\min}(\mathbf{X}_j^t) \geq \frac{1}{2}p \cdot \sigma_{\min}(\mathbf{R}_U)^2 \geq C \frac{p \cdot \sigma_{\min}(\mathbf{\Sigma}^*)^2}{\kappa^2} > 0, \quad (11)$$

where the last inequality follows from (8).

This shows that X used in update of $\widehat{\mathbf{V}}^{t+1}$ is PSD, and hence the update for $\widehat{\mathbf{V}}^{t+1}$ is given by:

$$\mathbf{R}_U(\widehat{\mathbf{V}}^{t+1})^\top \quad (12)$$

$$= (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{C}_j \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top + \bar{\mathbf{g}}_j^t) \quad (13)$$

$$= (\mathbf{U}^t)^\top \mathbf{U}^* \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top \quad (14)$$

$$- (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{D}_j (\mathbf{U}^t)^\top \mathbf{U}^* - \mathbf{C}_j) \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top - (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{N}_j^t (\mathbf{U}^t)^\top \mathbf{U}^* \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top - \bar{\mathbf{g}}_j^t), \quad (15)$$

where $\mathbf{D}_j = (\mathbf{U}^t)^\top \left(\frac{1}{p} \sum_{i:(i,j) \in \Omega} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{U}^t$, $\mathbf{C}_j = (\mathbf{U}^t)^\top \left(\frac{1}{p} \sum_{i:(i,j) \in \Omega} \mathbf{e}_i \mathbf{e}_i^\top \right) \mathbf{U}^*$, and $\bar{\mathbf{g}}_j^t = \frac{1}{p} \mathbf{R}_U^{-1} \mathbf{g}_j^t$.

That is,

$$\begin{aligned} \widehat{\mathbf{V}}^{t+1} \mathbf{R}_U &= \mathbf{V}^* \mathbf{\Sigma}^* (\mathbf{U}^*)^\top \mathbf{U}^t - \mathbf{E}^\top, \quad \mathbf{E}_j = \mathbf{E}_j^1 + \mathbf{E}_j^2, \\ \mathbf{E}_j^1 &= (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{D}_j (\mathbf{U}^t)^\top \mathbf{U}^* - \mathbf{C}_j) \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top, \quad \mathbf{E}_j^2 = (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{N}_j^t (\mathbf{U}^t)^\top \mathbf{U}^* \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top - \bar{\mathbf{g}}_j^t). \end{aligned} \quad (16)$$

Let $\widehat{\mathbf{V}}^{t+1} = \mathbf{V}^{t+1} \mathbf{R}_V$. Then,

$$\begin{aligned} \mathbf{V}^{t+1} \mathbf{R}_V \mathbf{R}_U &= \mathbf{V}^* \mathbf{\Sigma}^* (\mathbf{U}^*)^\top \mathbf{U}^t - \mathbf{E}^\top, \quad \mathbf{E}_j = \mathbf{E}_j^1 + \mathbf{E}_j^2, \\ \mathbf{E}_j^1 &= (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{D}_j (\mathbf{U}^t)^\top \mathbf{U}^* - \mathbf{C}_j) \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top, \quad \mathbf{E}_j^2 = (\mathbf{D}_j + \mathbf{N}_j^t)^{-1} (\mathbf{N}_j^t (\mathbf{U}^t)^\top \mathbf{U}^* \mathbf{\Sigma}^* (\mathbf{V}^*)_j^\top - \bar{\mathbf{g}}_j^t). \end{aligned} \quad (17)$$

Using the technique of (Jain et al., 2013, Lemma 5.6) and the bound on $\sigma_{\min}(\mathbf{D}_j + \mathbf{N}_j^t)$ (see (10)), we get:

$$\|(\mathbf{\Sigma}^*)^{-1} \mathbf{E}^1\|_F \leq \frac{C}{\kappa} \text{Err}(\mathbf{U}^t, \mathbf{U}^*). \quad (18)$$

Similarly, w.p. $\geq 1 - 1/n^{100}$:

$$\begin{aligned} \|(\mathbf{\Sigma}^*)^{-1} \mathbf{E}^2\|_F &\leq \frac{2}{\sigma_{\min}(\mathbf{\Sigma}^*)} \cdot \left(\frac{\Gamma_{\mathbf{u}}^2 \sigma}{p \sigma_{\min}(\mathbf{R}_U)^2} \frac{\sigma_{\max}(\mathbf{\Sigma}^*) \mu \sqrt{r^2 m \log n}}{\sqrt{m}} + \frac{\Gamma_{\mathbf{u}} \Gamma_{\mathbf{M}} \sigma}{p \sigma_{\min}(\mathbf{R}_U)} \cdot \sqrt{mr \log n} \right), \\ &\leq \frac{C \sigma \sqrt{\log n}}{pn} \cdot (\kappa^5 \cdot \mu^3 r^2 + \mu^3 r^2 \kappa^3) \leq \frac{C \kappa^5 \cdot \mu^3 r^2 \sqrt{\log n}}{\sqrt{pn}} \cdot \frac{\sqrt{m \log n} \cdot T \sqrt{\log(1/\delta)}}{\varepsilon}. \end{aligned} \quad (19)$$

Let $\beta = \frac{C \kappa^5 \cdot \mu^3 r^2 \sqrt{\log n}}{\sqrt{pn}} \frac{\sqrt{m \log n} \cdot T \log(1/\delta)}{\varepsilon}$. Now,

$$\sigma_{\min}(\mathbf{R}_V \mathbf{R}_U) \geq \sigma_{\min}(\mathbf{\Sigma}^*) (1 - 2\text{Err}(\mathbf{U}^t, \mathbf{U}^*) - \kappa \beta) \geq \frac{\sigma_{\min}(\mathbf{\Sigma}^*)}{2}, \quad (20)$$

where the last inequality holds because:

$$\sqrt{pn} \geq C \kappa^6 \mu^3 r^2 \sqrt{m} \frac{T \log n \sqrt{\log(1/\delta)}}{\varepsilon}.$$

Using (17), we have:

$$\max_j \|(\mathbf{V}^{t+1})_j^\top\|_2 \leq \frac{2\mu\kappa\sqrt{r}}{\sqrt{m}} + \frac{4\mu\kappa\sqrt{r}}{\sqrt{m}} + \frac{2\mu\kappa r \Gamma_{\mathbf{u}}^2 \sigma}{\sqrt{m} p \sigma_{\min}(\mathbf{R}_U)^2} + \frac{2\Gamma_{\mathbf{u}} \Gamma_{\mathbf{M}} \sigma \sqrt{r}}{p \sigma_{\min}(\mathbf{\Sigma}^*) \sigma_{\min}(\mathbf{R}_U)} \leq \frac{16\mu\kappa\sqrt{r}}{\sqrt{m}}, \quad (21)$$

where the last inequality follows from the assumption that $\sqrt{pn} \geq C \kappa^6 \mu^3 r^2 \sqrt{m} \frac{T \log n \sqrt{\log(1/\delta)}}{\varepsilon}$. This concludes the proof.

B.2.2. PROOF OF LEMMA 4

The proof for this key Lemma follows technique similar to the above proof. That is, using previous lemma, the clipping operations do not have any effect, and hence we get noisy ALS updates. Now, using (18), (19), (20), and Lemma 8, we have:

$$\text{Err}(\mathbf{V}^*, \mathbf{V}^{t+1}) \leq \frac{1}{4} \text{Err}(\mathbf{U}^*, \mathbf{U}^t) + 4\kappa\beta. \quad (22)$$

This proves the lemma.

B.3. Proof of Theorem 2

Using Lemma 8, we have:

$$\text{Err}(\mathbf{V}^*, \mathbf{V}^t) \leq \frac{1}{4} + \text{Err}(\mathbf{V}^*, \mathbf{V}^{t+1}) + \alpha,$$

where $\alpha \leq \frac{C\kappa^6 \cdot \mu^3 r^2 \sqrt{\log n} \sqrt{m \log n} \cdot T \log 1/\delta}{\sqrt{pn} \varepsilon}$. So, after $T = \log \frac{\text{Err}(\mathbf{V}^*, \mathbf{V}^0)}{\alpha}$ iterations, $\text{Err}(\mathbf{V}^*, \mathbf{V}^T) \leq 2\alpha$.

As $\hat{\mathbf{U}}^T = \arg \min_{\hat{\mathbf{U}}} \|\mathbf{M} - \hat{\mathbf{U}}^T (\mathbf{V}^T)^\top\|_F$, we have:

$$\|\mathbf{M} - \hat{\mathbf{U}}^T (\mathbf{V}^T)^\top\|_F \leq \|\mathbf{M} - \mathbf{U}^* \boldsymbol{\Sigma}^* (\mathbf{V}^T)^\top\|_F \leq \|\mathbf{M}\|_F \|\mathbf{V}^* - \mathbf{V}^T\|_2 \leq 2\alpha \|\mathbf{M}\|_F, \quad (23)$$

where last inequality follows from the fact that $\|\mathbf{V}^* - \mathbf{V}^T\|_2 \leq 2\text{Err}(\mathbf{V}^*, \mathbf{V})$.

This shows the second claim of the theorem. The third claim follows similarly while using incoherence of \mathbf{V}^T .

Lemma 8. Let $\hat{\mathbf{U}} = \mathbf{U}^* \boldsymbol{\Sigma}^* \mathbf{W} + \mathbf{E}$ and $\mathbf{U} = \hat{\mathbf{U}} \mathbf{R}^{-1}$ where $\boldsymbol{\Sigma}^*$ is a diagonal matrix, $\mathbf{W} \in \mathbb{R}^{r \times r}$, and $\mathbf{R}^2 = \hat{\mathbf{U}}^\top \hat{\mathbf{U}}$. Then, assuming $\sigma_{\min}(\boldsymbol{\Sigma}^*) \sigma_{\min}(\mathbf{W}) > \|\boldsymbol{\Sigma}^*\|_2 \|\mathbf{E}(\boldsymbol{\Sigma}^*)^{-1}\|_2$, the following holds:

$$\|(\mathbf{I} - \mathbf{U}^* (\mathbf{U}^*)^\top) \mathbf{U}\|_2 \leq \frac{\|\mathbf{E} \cdot (\boldsymbol{\Sigma}^*)^{-1}\|_2}{\frac{\sigma_{\min}(\boldsymbol{\Sigma}^*)}{\|\boldsymbol{\Sigma}^*\|_2} \sigma_{\min}(\mathbf{W}) - \|\mathbf{E}(\boldsymbol{\Sigma}^*)^{-1}\|_2}.$$

That is,

$$\|(\mathbf{I} - \mathbf{U}^* (\mathbf{U}^*)^\top) \mathbf{U}\|_2 \leq \frac{\kappa \|\mathbf{E}\|_2}{\sigma_{\min}(\boldsymbol{\Sigma}^*) \sigma_{\min}(\mathbf{W}) - \kappa \|\mathbf{E}\|_2}.$$

Proof.

$$\|(\mathbf{I} - \mathbf{U}^* (\mathbf{U}^*)^\top) \mathbf{U}\|_2 \leq \|\mathbf{E} \cdot \mathbf{R}^{-1}\|_2 \leq \|\mathbf{E}(\boldsymbol{\Sigma}^*)^{-1}\|_2 \|\boldsymbol{\Sigma}^* \mathbf{R}^{-1}\|_2.$$

Furthermore, $\|\boldsymbol{\Sigma}^* \mathbf{R}^{-1}\|_2 \leq \|\boldsymbol{\Sigma}^*\|_2 \|\mathbf{R}^{-1}\|_2$. Now,

$$\frac{1}{\|\mathbf{R}^{-1}\|_2} = \sigma_{\min}(\mathbf{R}) \geq \sigma_{\min}(\boldsymbol{\Sigma}^*) \sigma_{\min}(\mathbf{W}) - \|\boldsymbol{\Sigma}^*\|_2 \|\mathbf{E}(\boldsymbol{\Sigma}^*)^{-1}\|_2.$$

That is,

$$\|(\mathbf{I} - \mathbf{U}^* (\mathbf{U}^*)^\top) \mathbf{U}\|_2 \leq \frac{\|\mathbf{E} \cdot (\boldsymbol{\Sigma}^*)^{-1}\|_2}{\frac{\sigma_{\min}(\boldsymbol{\Sigma}^*)}{\|\boldsymbol{\Sigma}^*\|_2} \sigma_{\min}(\mathbf{W}) - \|\mathbf{E} \cdot (\boldsymbol{\Sigma}^*)^{-1}\|_2}.$$

□

B.4. Noisy Power Iteration Initialization

In this section, we derive a tighter initialization routine through the noisy power iteration procedure. We will show that it only requires $n = \tilde{O}(m)$ and can succeed with high probability.

Prior work on noisy power iteration requires $\mathbf{A} = \mathbf{P}_\Omega(\mathbf{M})^\top \mathbf{P}_\Omega(\mathbf{M})$ to be incoherent (Hardt & Roth, 2012; 2013; Hardt & Price, 2013). But whether this is true under our sparsity condition is a difficult problem. For example, related bounds, first

conjectured in (Dekel et al., 2011), have only been shown to hold for constant p (Vu & Wang, 2015; Rudelson & Vershynin, 2015) but still open for $p = O(\log^c m/m)$ for constant $c > 0$, the range interesting to us. To overcome this difficulty, we show that we actually do not need \mathbf{A} to be fully incoherent. Instead, we just need \mathbf{A} 's top- r eigenspaces to be incoherent, and the existence of a (moderate) gap between the top eigenspaces and the rest, both of which we are able to establish. Given these two conditions, we then add the proper amount of noise, with a magnitude in between the top- r eigenvalues and the rest, such that i) it does not interfere with the ‘‘boosting’’ of the top- r eigenspace; and ii) it ‘‘randomizes’’ the remaining eigenvectors such that their incoherence is preserved through the power iteration.

For simplicity, we will present a detailed proof in the rank-1 case. Here, we say a vector $\mathbf{w} \in \mathbb{R}^m$ is μ -incoherent if $\|\mathbf{w}\|_\infty \leq \frac{\mu}{\sqrt{m}}$.

Algorithm 2 Noisy power iteration.

Required: $P_\Omega(\mathbf{M}) \in \mathbb{R}^{n \times m}$, number of iterations: T , incoherence parameter: ν , s : threshold for maximum number of ratings per user, entry clipping parameter: Γ_M .

- 1 $\mathbf{w}_1 \leftarrow$ Random unit vector in m -dimensions.
 - for** $1 \leq t \leq T$ **do**
 - 2 **If** \mathbf{w}_t is not ν -incoherent, **then** report failure and stop.
 - 3 Compute $\mathbf{z}_t = (P_\Omega(\mathbf{M})^\top P_\Omega(\mathbf{M})) \cdot \mathbf{w}_t$, and $\tilde{\mathbf{z}} \leftarrow \mathbf{z}_t + \mathbf{g}_t$, where $\mathbf{g}_t \sim \mathcal{N}(0, \sigma^2 \cdot \mathbb{I})$.
 - 4 Normalize $\tilde{\mathbf{z}}_t$ to obtain \mathbf{w}_{t+1} .
 - end**
 - return** \mathbf{w}_{T+1} .
-

Theorem 9 (Privacy guarantee). *Algorithm 2 satisfies $(\alpha, \alpha\rho^2)$ -RDP, where $\rho^2 = \frac{Ts^3\Gamma_M^4\nu^2}{2m\sigma^2}$.*

The proof follows immediately from ℓ_2 -sensitivity analysis and the RDP guarantee for Gaussian mechanism (Mironov, 2017).

For the utility guarantee, consider the case where Ω is randomly sampled with probability p (so setting $s \approx mp$ is sufficient). Furthermore, since we are in the rank-1 case, we will assume $\mathbf{M} = \sqrt{mn} \cdot \mathbf{u} \otimes \mathbf{v}$, where both \mathbf{u} and \mathbf{v} are μ -incoherent so Γ_M can be set as μ^2 . Below we assume $\mu, \Gamma_M = O(1)$ for simplicity. Now we will show that we can set $p = O(\log^3 m/m)$ and $n = O(m \log m \sqrt{\log(1/\delta)}/\varepsilon)$, with proper choices of ν, σ , such that $\rho^2 \leq (\varepsilon + \log(1/\delta))/\varepsilon^2$, and the above procedure returns a vector \mathbf{w} such that $|\mathbf{w} \cdot \mathbf{v}| > 0.6$ with probability $1 - o(1)$, which can be boosted to high probability by the standard method.

Theorem 10 (Utility guarantee). *There exists constant $C_1, C_2 > 0$, such that for any $\delta \in (0, 1), \varepsilon \in (0, \log(1/\delta))$, if $p \geq \frac{C_1 \log^3 m}{m}$ and $\sqrt{pn} \geq C_2 \frac{\sqrt{\log(1/\delta)}}{\varepsilon} \cdot \sqrt{m} \log^{5/2} m$, then we can choose settings of incoherence parameter ν , noise standard deviation σ , and number of time steps T in Algorithm 2 s.t., w.p. $1 - o(1)$, we have $|\mathbf{w}_{T+1} \cdot \mathbf{v}| > 0.6$, where \mathbf{v} is the right singular vector of \mathbf{M} , and $\rho = \frac{\varepsilon}{2\sqrt{\log(1/\delta)}}$, i.e. Algorithm 2 satisfies (ε, δ) -differential privacy. The probability guarantee can be boosted to high probability $1 - m^{-c}$ for any $c > 0$ with the private selection algorithm (Liu & Talwar, 2019; Zhu & Wang, 2020).*

Proof. We will prove this theorem through a sequence of claims. Write $\lambda = p^2 mn$. The following claim is from previous work, e.g. (Recht, 2011).

Claim 11. *For $n = \Omega(m), p = \Omega(\frac{\log m}{m})$, with high probability, $\mathbf{A} = \lambda(\mathbf{v} \otimes \mathbf{v}) + \mathbf{B}$, where $\|\mathbf{B}\|_2 = O(np\sqrt{\log m})$. Hence, if λ_1, \mathbf{h}_1 are the principal eigenvalue and eigenvector, respectively, of \mathbf{A} , then $\lambda_1 = (1 \pm o(1))\lambda$ and $|\mathbf{h}_1 \cdot \mathbf{v}| = 1 - o(1)$.*

One key fact we need is that \mathbf{h}_1 is not only close to \mathbf{v} , but also incoherent. The proof essentially follows the arguments in the proof of Theorem 2.16 in (Erdős et al., 2013). One difference in our case is that the entries in \mathbf{A} are not independent because $\mathbf{A} = P_\Omega(\mathbf{M})^\top P_\Omega(\mathbf{M})$. This difficulty was overcome in (Jain & Netrapalli, 2015)(Lemma 6) by using the resampling technique. Here we present a direct argument, which might be of independent interest.

Claim 12. *For $n = \Omega(m), p = \Omega(\frac{\log^2 m}{m})$, with high probability, the principal eigenvector of \mathbf{A} is C -incoherent for some absolute constant $C > 0$.*

Proof. We treat $P_\Omega(\mathbf{M})$ as the adjacency matrix of a random bipartite graph and apply the techniques similar to (Erdős et al., 2013; Jain & Netrapalli, 2015). Recall \mathbf{h}_1 is the principal eigenvector of \mathbf{A} with the eigenvalue λ_1 . Denote by $\mathbf{N} = P_\Omega(\mathbf{M}) - p\sqrt{mn}(\mathbf{u} \otimes \mathbf{v})$. We will first show that when $np \gg \log m$, with high probability, there exists \mathbf{w} , where $\|\mathbf{w}\|_\infty = O(1/\sqrt{m})$, such that

$$\mathbf{h}_1 = (1 \pm o(1)) \left(\mathbb{I} - \frac{1}{\lambda_1} \mathbf{N}^\top \mathbf{N} \right)^{-1} \mathbf{w}. \quad (24)$$

Since $\lambda_1 = (1 \pm o(1))\lambda$, $\|\mathbf{N}^\top \mathbf{N}\|_2 = o(\lambda_1)$, we can expand the above equation to:

$$\mathbf{h}_1 = (1 \pm o(1)) \sum_{k \geq 0} \left(\frac{1}{\lambda_1} \mathbf{N}^\top \mathbf{N} \right)^k \mathbf{w}.$$

We then apply the method in (Erdős et al., 2013; Jain & Netrapalli, 2015) to show that there exists constant $c_0 < 1$ such that with high probability:

$$\left\| \left(\frac{1}{\lambda_1} \mathbf{N}^\top \mathbf{N} \right)^k \mathbf{w} \right\|_\infty \leq c_0^k \|\mathbf{w}\|_\infty. \quad (25)$$

These would imply that \mathbf{h}_1 is $O(1)$ -incoherent. We first prove (24). Since $\mathbf{N} = P_\Omega(\mathbf{M}) - p\sqrt{mn}(\mathbf{u} \otimes \mathbf{v})$, we can write

$$\mathbf{A} = P_\Omega(\mathbf{M})^\top P_\Omega(\mathbf{M}) = \mathbf{N}^\top \mathbf{N} + p\sqrt{mn}(\mathbf{v} \otimes \mathbf{u})P_\Omega(\mathbf{M}) + p\sqrt{mn}P_\Omega(\mathbf{M})^\top(\mathbf{u} \otimes \mathbf{v}) - p^2mn(\mathbf{v} \otimes \mathbf{v}).$$

We first observe that with high probability $\tilde{\mathbf{v}} = \frac{1}{p\sqrt{mn}}P_\Omega(\mathbf{M})^\top \mathbf{u}$ satisfies that

$$|\tilde{v}_j - v_j| = O\left(\sqrt{\frac{\log m}{pn}}|v_j|\right) \quad \text{for all } j \in [m]. \quad (26)$$

Since $|\mathbf{h}_1 \cdot \mathbf{v}| = 1 - o(1)$, we have that when $pn \gg \log m$, $|\mathbf{h}_1 \cdot \tilde{\mathbf{v}}_j - \mathbf{h}_1 \cdot \mathbf{v}| = o(1)$. Hence

$$p\sqrt{mn}(\mathbf{v} \otimes \mathbf{u})P_\Omega(\mathbf{M})\mathbf{h}_1 = p\sqrt{mn}(P_\Omega(\mathbf{M})^\top \mathbf{u} \cdot \mathbf{h}_1)\mathbf{v} = p^2mn((\mathbf{h}_1 \cdot \mathbf{v}) \pm o(1))\mathbf{v}. \quad (27)$$

In addition,

$$p\sqrt{mn}P_\Omega(\mathbf{M})^\top(\mathbf{u} \otimes \mathbf{v})\mathbf{h}_1 = p^2mn(\mathbf{h}_1 \cdot \mathbf{v})\tilde{\mathbf{v}}. \quad (28)$$

Applying (27) and (28), we have that

$$\begin{aligned} \mathbf{A}\mathbf{h}_1 &= (\mathbf{N}^\top \mathbf{N} + p\sqrt{mn}(\mathbf{v} \otimes \mathbf{u})P_\Omega(\mathbf{M}) + p\sqrt{mn}P_\Omega(\mathbf{M})^\top(\mathbf{u} \otimes \mathbf{v}) - p^2mn(\mathbf{v} \otimes \mathbf{v}))\mathbf{h}_1 \\ &= \mathbf{N}^\top \mathbf{N}\mathbf{h}_1 + p^2mn((\mathbf{h}_1 \cdot \mathbf{v}) \pm o(1))\mathbf{v} + p^2mn(\mathbf{h}_1 \cdot \mathbf{v})\tilde{\mathbf{v}} - p^2mn(\mathbf{h}_1 \cdot \mathbf{v})\mathbf{v} \\ &= \mathbf{N}^\top \mathbf{N}\mathbf{h}_1 + \lambda(\mathbf{h}_1 \cdot \mathbf{v})(\tilde{\mathbf{v}} \pm o(1)\mathbf{v}). \end{aligned}$$

Let $\mathbf{w} = (\mathbf{h}_1 \cdot \mathbf{v})(\tilde{\mathbf{v}} \pm o(1)\mathbf{v})$. Recall $\|\mathbf{v}\|_\infty = O(1/\sqrt{m})$. When $pn \gg \log m$, using (26), we have $\|\tilde{\mathbf{v}}\|_\infty = O(1/\sqrt{m})$ too. Hence $\|\mathbf{w}\|_\infty = O(1/\sqrt{m})$. Since $\mathbf{A}\mathbf{h}_1 = \lambda_1\mathbf{h}_1$, we have that $\mathbf{N}^\top \mathbf{N}\mathbf{h}_1 + \lambda\mathbf{w} = \lambda_1\mathbf{h}_1$, hence $(\lambda_1\mathbb{I} - \mathbf{N}^\top \mathbf{N})\mathbf{h}_1 = \lambda\mathbf{w}$, which implies (24).

Now we prove (25). Since we $\lambda_1 = (1 \pm o(1))\lambda$, it suffices to prove (25) by replacing λ_1 with λ instead. Furthermore, since $\|\mathbf{N}^\top \mathbf{N}\|_2 = o(\lambda)$, it suffices to consider $k = O(\log m)$. Let $\mathbf{N}' = \frac{1}{p\sqrt{mn}}\mathbf{N}$. Write $\mathbf{w}' = \left(\frac{1}{\lambda}\mathbf{N}^\top \mathbf{N}\right)^k \mathbf{w} = (\mathbf{N}'^\top \mathbf{N}')^k \mathbf{w}$. Following the proof of Lemma 7.10 in (Erdős et al., 2013), we will bound the q -th moment of $|w'_j|$ and apply the Markov inequality. We treat \mathbf{N}' as the adjacency matrix of a bipartite graph $[n] \times [m]$ where each edge is labeled with a random

variable $\xi_{ij} = (\frac{1}{p}\chi_{ij} - 1)u_iv_j$ where χ_{ij} 's are independent Bernoulli random variables taking value 1 with probability p . By its definition $\mathbf{E}[\xi_{ij}] = 0$, and for $r \geq 2$,

$$\mathbf{E}[|\xi_{ij}|^r] = p\left(\frac{1}{p} - 1\right)|u_iv_j|^r + (1-p)(|u_iv_j|)^r = O\left(p\frac{1}{(p\sqrt{mn})^r}\right). \quad (29)$$

Let $G = [n] \times [m]$ denote the complete bipartite graph with each edge ij labeled with ξ_{ij} . For $j_1, j_2 \in [m]$, let $\mathcal{P}_k(j_1, j_2)$ denote all the length $2k$ paths in G starting from the node j_1 and ending at node j_2 . Then,

$$w'_j = \sum_{j_1} w_{j_1} \sum_{P \in \mathcal{P}_k(j_1, j)} \prod_{e \in P} \xi_e.$$

And

$$w_j'^q = \sum_{j_1, j_2, \dots, j_q} w_{j_1} w_{j_2} \dots w_{j_q} \sum_{\forall \ell P_\ell \in \mathcal{P}_k(j_\ell, j)} \prod_{e \in \cup_\ell P_\ell} \xi_e.$$

Here $\cup_\ell P_\ell$ is understood as a multiple set. Hence

$$\begin{aligned} \mathbf{E}[w_j'^q] &= \sum_{j_1, j_2, \dots, j_q} w_{j_1} w_{j_2} \dots w_{j_q} \sum_{\forall \ell P_\ell \in \mathcal{P}_k(j_\ell, j)} \mathbf{E}\left[\prod_{e \in \cup_\ell P_\ell} \xi_e\right] \\ &\leq \|w\|_\infty^q \sum_{j_1, j_2, \dots, j_q} \sum_{\forall \ell P_\ell \in \mathcal{P}_k(j_\ell, j)} |\mathbf{E}\left[\prod_{e \in \cup_\ell P_\ell} \xi_e\right]|. \end{aligned}$$

We will now bound

$$\sum_{j_1, j_2, \dots, j_q} \sum_{\forall \ell P_\ell \in \mathcal{P}_k(j_\ell, j)} |\mathbf{E}\left[\prod_{e \in \cup_\ell P_\ell} \xi_e\right]|. \quad (30)$$

Since $\mathbf{E}[\xi_{ij}] = 0$ and ξ_{ij} 's are independent random variables, for $\mathbf{E}\left[\prod_{e \in \cup_\ell P_\ell} \xi_e\right] \neq 0$, it must be that every edge in $\cup_\ell P_\ell$ appears at least twice. The bound on $\mathbf{E}[w_j'^q]$ is by counting the number of such paths and applying the moments bound (29).

The argument follows (Erdős et al., 2013). We give it here for completeness. Let P denote the set of edges, without multiplicity, in $\cup_\ell P_\ell$. Write $t = |P|$. Then $t \leq kq$ since every edge in $\cup_\ell P_\ell$ has to appear at least twice. In addition, the edges in P form a connected component because all the paths are connected to node j . Hence there are at most $t + 1$ vertices. Since the set must include j , there are at most $\binom{m+n}{t}$ choices of the set of vertices. Among $t + 1$ vertices, we can make at most $\binom{t+2k}{2k} (2k)!$ different paths of length $2k$. Hence the total number of paths is bounded by

$$\binom{m+n}{t} \left(\binom{t+2k}{2k} (2k)! \right)^q \leq n^t (c_1 kq)^{2kq},$$

for some $c_1 > 0$. In the last inequality, we used that $n = \Omega(m)$ and $t \leq kq$.

Suppose that $P = \{e_1, e_2, \dots, e_t\}$, and the multiplicity of these edges in $\cup_\ell P_\ell$ are s_1, s_2, \dots, s_t respectively. So $\sum_i s_i = 2kq$. Using (29),

$$\begin{aligned} \mathbf{E}\left[\prod_{e \in \cup_\ell P_\ell} |\xi_e|\right] &= \mathbf{E}[|\xi_{e_1}|^{s_1}] \mathbf{E}[|\xi_{e_2}|^{s_2}] \dots \mathbf{E}[|\xi_{e_t}|^{s_t}] \\ &= \left(p\frac{1}{(p\sqrt{mn})^{s_1}}\right) \dots \left(p\frac{1}{(p\sqrt{mn})^{s_t}}\right) \\ &= p^t \frac{1}{(p\sqrt{mn})^{2kq}}. \end{aligned}$$

Hence the contribution to (30) by the case of $|P| = t$ is bounded by:

$$n^t (c_1 kq)^{2kq} p^t \frac{1}{(p\sqrt{mn})^{2kq}} = \frac{(c_1 kq)^{2kq}}{(pm)^{kq} (pn)^{kq-t}}.$$

Since $t \leq kq$, (30) is bounded by

$$O\left(\frac{(c_1 k q)^{2kq}}{(pm)^{kq}}\right).$$

Fix $c_0 = 1/2$. For any $c > 0$ and $k = O(\log m)$, we can choose $q = \log m / (2k)$ and $p \geq c' \log^2 m / m$ for some $c' > 0$ such that (30) is bounded by $c_0^{kq} m^{-c}$. Applying Markov inequality, we have that with probability $1 - m^{-c}$, $|w'_j| = O(c_0^k \|w\|_\infty)$. Since c can be chosen arbitrarily, we have that with high probability this holds for all $k = O(\log m)$. This completes the proof. \square

The following summarizes the above claims and the conditions we need in our proof.

Claim 13. Assume $n = \Omega(m)$, $p = \Omega(\frac{\log^3 m}{m})$. Let $\mathbf{A} = \sum_i \lambda_i (\mathbf{h}_i \otimes \mathbf{h}_i)$ be the eigen-decomposition of \mathbf{A} where $\lambda_1 \geq \lambda_2 \geq \dots \lambda_m \geq 0$. Then $\lambda_1 = (1 \pm o(1)) p^2 m n$, and for $i \geq 2$, $\lambda_i = O(np \sqrt{\log m}) = o(\lambda_1 / \log^2 m)$. In addition $\mathbf{h}_1 \cdot \mathbf{v} = 1 - o(1)$, and $\|\mathbf{h}_1\|_\infty \leq C/\sqrt{m}$.

Now we show that

Claim 14. There exists $c_2, c_3 > 0$, such that if $\sigma \geq c_2 \frac{\lambda}{\sqrt{m} \log^{3/2} m}$, then $\forall t \in [T]$, \mathbf{w}_t 's are $c_3 \sqrt{\log m}$ -incoherent.

Proof. For notation convenience, we set $\mathbf{w}_0 = 0$, hence \mathbf{w}_1 is a random unit vector. Write $\mathbf{w}_t = \sum_i \alpha_{ti} \mathbf{h}_i$. Then $\mathbf{A} \mathbf{w}_t = \sum_i \lambda_i \alpha_{ti} \mathbf{h}_i$ and $\tilde{\mathbf{z}}_t = \mathbf{A} \mathbf{w}_t + \mathbf{g}_t = \sum_i (\lambda_i \alpha_{ti} + g_{ti}) \mathbf{h}_i$, where $g_{ti} = g_t \cdot \mathbf{h}_i$. Hence $\alpha_{(t+1)i} = (\lambda_i \alpha_{ti} + g_{ti}) / \|\tilde{\mathbf{z}}_t\|_2$. Since \mathbf{g}_t is sampled from $\mathcal{N}(0, \sigma^2 \mathbb{I})$, g_{ti} 's are i.i.d Gaussian from $\mathcal{N}(0, \sigma^2)$. In addition $\alpha_{0i} = 0$. By induction, we have that $\mathbf{E}[\alpha_{ti}] = 0$, and the signs $\text{sign}(\alpha_{t1}), \dots, \text{sign}(\alpha_{tm})$ are uniformly distributed in $\{-1, 1\}^m$, independent of their values, by the same argument in (Hardt & Price, 2013) Lemma 4.13. We will now first show, by induction, that with high probability, $\max_{i \geq 2} |\lambda_i \alpha_{ti}| = O(\sigma)$.

When $t = 0$, this is clearly true. Now suppose this holds for w_t . All the following statements hold with high probability. Note in the following α_{ti} 's are random variables. By $\|\tilde{\mathbf{z}}_t\|^2 = \sum_i (\lambda_i \alpha_{ti} + g_{ti})^2$ and $\max_{i \geq 2} |\lambda_i \alpha_{ti}| = O(\sigma)$ with high probability, we have that with high probability,

$$\|\tilde{\mathbf{z}}_t\|_2^2 = \Omega(\lambda_1^2 \alpha_{t1}^2 + m \sigma^2). \quad (31)$$

By induction hypothesis for any $i \geq 2$, $|\lambda_i \alpha_{ti}| = O(\sigma)$, hence $|\lambda_i \alpha_{ti} + g_{ti}| = O(\sqrt{\log m} \sigma)$. Since $\|\tilde{\mathbf{z}}_t\| = \Omega(\sqrt{m} \sigma)$, $|\alpha_{(t+1),i}| = |\lambda_i \alpha_{ti} + g_{ti}| / \|\tilde{\mathbf{z}}_t\| = O(\frac{\sqrt{\log m}}{\sqrt{m}})$, hence $|\lambda_i \alpha_{(t+1),i}| = O(\frac{\lambda}{\log^2 m} \cdot \frac{\sqrt{\log m}}{\sqrt{m}}) = O(\frac{\lambda}{\log^{3/2} m \sqrt{m}}) = O(\sigma)$. We can clearly choose c_3 large enough to make sure the induction goes through.

The j -th coordinate of $\tilde{\mathbf{z}}_t$ is $\tilde{\mathbf{z}}_t(j) = \sum_i (\lambda_i \alpha_{ti} + g_{ti}) h_i(j)$. Consider $\tilde{\mathbf{z}}_t(j)' = \sum_{i \geq 2} (\lambda_i \alpha_{ti} + g_{ti}) h_i(j)$. Clearly g_{ti} 's are independent Gaussian variables. In addition, $\text{sign}(\alpha_{t1}), \dots, \text{sign}(\alpha_{tm})$ are uniformly distributed over $\{-1, 1\}^m$. Since $|\lambda_i \alpha_{ti}| = O(\sigma)$ for $i \geq 2$, and $\sum_i h_i(j)^2 \leq 1$, by applying the concentration bound, we have that with high probability

$$|\tilde{\mathbf{z}}_t(j)'| = O\left(\sqrt{\log m \sum_{i \geq 2} \sigma^2 h_i(j)^2}\right) = O(\sigma \sqrt{\log m}).$$

Hence, with high probability,

$$\tilde{\mathbf{z}}_t(j)^2 = O(\lambda_1^2 \alpha_{t1}^2 h_1(j)^2 + (\log m) \sigma^2) = O(\lambda_1^2 \alpha_{t1}^2 / m + (\log m) \sigma^2). \quad (32)$$

The last inequality is by $\|\mathbf{h}_1\|_\infty \leq C/\sqrt{m}$. Combining with Equation (31), by distinguishing the cases of $|\lambda_1 \alpha_{t1}| \leq \sqrt{m} \sigma$ and $|\lambda_1 \alpha_{t1}| \geq \sqrt{m} \sigma$, we have that $\frac{|\tilde{\mathbf{z}}_t(j)|}{\|\tilde{\mathbf{z}}_t\|_2} = O(\frac{\sqrt{\log m}}{\sqrt{m}})$, i.e w_t is $O(\sqrt{\log m})$ -incoherent. \square

Now, we show that we need only $O(\log m)$ round to get a constant approximation to v . Note that here we cannot get high probability bound because we need the initial $|\alpha_{01}|$ to be $\Omega(1/\sqrt{m \log m})$ to bootstrap the process. But it does happen with probability $1 - o(1)$.

Claim 15. *There exists $c_4, c_5 > 0$, such that with n, p, σ, ν as set above, if $T \geq c_4 \log m$, $|w_T \cdot v| \geq 1 - c_5/\log m$ w.p. $1 - o(1)$.*

Proof. It suffices to show that $|\alpha_{T1}| = \Omega(1)$. Because $\lambda_2 = O(\lambda_1/\log^2 m)$ and $\sqrt{m \log m} \sigma = O(\lambda_1/\log m)$, once $|\alpha_{T1}| = \Omega(1)$, with one more round we would have $|\alpha_{T+1,1}| = 1 - O(1/\log m)$.

Suppose that $|\alpha_{01}| = \Omega(1/\sqrt{m \log m})$, which happens with probability $1 - o(1)$. Then $|\lambda_1 \alpha_{01}| = \Omega(\log m) \sigma$. We can prove by induction that $|\lambda_1 \alpha_{t1}| = \Omega(\log m) \sigma$ with high probability. By $\alpha_{(t+1)1} = (\lambda_1 \alpha_{t1} + g_{t1})/\|\tilde{z}_t\|_2$, we have $|\alpha_{(t+1)1}| \geq \frac{1}{2} |\lambda_1 \alpha_{t1}|/\|\tilde{z}_t\|_2$ with high probability. By Equation (32), we can bound

$$\|\tilde{z}_t\|_2^2 = O(\lambda_1^2 \alpha_{t1}^2 + (m \log m) \sigma^2).$$

Hence,

$$|\alpha_{(t+1)1}| \geq \frac{1}{2} |\lambda_1 \alpha_{t1}|/\|\tilde{z}_t\|_2 = \Omega(|\lambda_1 \alpha_{t1}|/(|\lambda_1 \alpha_{t1}| + \sqrt{m \log m} \sigma)). \quad (33)$$

Now if $|\lambda_1 \alpha_{t1}| \geq \sqrt{m \log m} \cdot \sigma$, then $|\alpha_{(t+1)1}| = \Omega(1)$ so we are done. Otherwise, $|\alpha_{(t+1)1}| = \Omega(|\lambda_1 \alpha_{t1}|/(\sqrt{m \log m} \cdot \sigma)) = \Omega(\log m) |\alpha_{t1}|$, by $\sigma = O\left(\frac{\lambda}{\sqrt{m \log^{3/2} m}}\right)$. Hence within $O(\log m)$ rounds, $\alpha_{t1} \geq c$ for some constant $c > 0$. \square

Now, gather the assumption $\Gamma_M = \mu^2 = O(1)$ and the conditions $\nu = O(\sqrt{\log m})$, $\sigma = \Omega\left(\frac{p^2 mn}{\sqrt{m \log^{3/2} m}}\right)$, $T = O(\log m)$ and plug them into the formula in Theorem 9. We have that

$$\rho^2 = \frac{T(mp)^3 \Gamma_M^4 \nu^2}{2m\sigma^2} = O\left(\frac{m \log^5 m}{pn^2}\right). \quad (34)$$

Hence we can set $n = O\left(\frac{\sqrt{\log(1/\delta)}}{\varepsilon} \sqrt{m/p} \log^{5/2} m\right)$, such that $\rho^2 = \frac{\varepsilon^2}{4 \log(1/\delta)}$. This completes the proof. We note that when $p = O(\log^3 m/m)$, $n = \tilde{O}(m \log m)$ which is nearly optimal.

Note that the only reason that prevents the high probability guarantee is due to the choice of the initial random vector. It can be boosted to high probability guarantee by running the process $O(\log m)$ times and privately releasing the vector w with $\|\mathbf{P}_\Omega(\mathbf{M})w\|^2$ above a threshold. This can be done through the private selection algorithm (Liu & Talwar, 2019) with an extra constant factor in ε . \square

C. Additional Details on Heuristic Improvements

Algorithm 3 summarizes the data pre-processing and sampling heuristics described in Section 5.

Algorithm 3 Data pre-processing heuristics.

Required: $P_\Omega(M)$: Observed ratings, Γ_M : entry clipping parameter, k : maximum number of ratings per user, σ_p : standard deviation of the pre-processing noise, β : fraction of movies to train on.

- 1 Clip entries in $P_\Omega(M)$ so that $\|P_\Omega(M)\|_\infty \leq \Gamma_M$
- 2 Uniformly sample Ω' :
 - for** $1 \leq i \leq n$ **do**
 - | $\Omega'_i \leftarrow$ sample k items from Ω_i uniformly.
 - end**
- 3 Compute movie counts $\tilde{c} \leftarrow \text{Counts}(\Omega')$.
- 4 Partition movies:
 - Let *Frequent* be the $\lceil \beta m \rceil$ movies with the largest \tilde{c} , and let *Infrequent* be the rest.
- 5 Adaptively sample Ω'' :
 - for** $1 \leq i \leq n$ **do**
 - | $\Omega''_i \leftarrow$ the k items in $(\Omega_i \cap \text{Frequent})$ with the lowest count \tilde{c} .
 - end**
- 6 Recompute movie counts $\tilde{c} \leftarrow \text{Counts}(\Omega'')$
- 7 Center the data $P_{\Omega''}(M) \leftarrow P_{\Omega''}(M) - \tilde{m}$, where $\tilde{m} = \frac{\sum_{(i,j) \in \Omega''} M_{ij} + \mathcal{N}(0, k\Gamma_M^2 \sigma_p^2)}{|\Omega''| + \mathcal{N}(0, k\sigma_p^2)}$

return $P_{\Omega''}(M), \tilde{c}$

Procedure *Counts*(Ω)

- for** $1 \leq j \leq m$ **do**
- | $\tilde{c}_j \leftarrow |\Omega_j| + \mathcal{N}(0, \sigma_p^2)$
- | **return** \tilde{c}
- end**

First, we compute differentially private movie counts (Line 3) using a uniform sample Ω' , and use it to partition the movies (Line 4) and to perform adaptive sampling (Line 5). The final subset used for training is Ω'' , which consists only of Frequent movies. Finally, to have a more accurate estimate of the counts, we recompute \tilde{c} on Ω'' (Line 6). We redo this computation as the counts are also used during optimization, as described in the next section. Note that in both computations of \tilde{c} , we use a subset of Ω that contains at most k movies per user, in order to guarantee user-level differential privacy.

Privacy accounting. As we saw in Theorem 1, Algorithm 1 with random initialization satisfies $(\alpha, \frac{\alpha(kT)}{2\sigma^2})$ -joint RDP. The data processing heuristics in Algorithm 3 satisfy $(\alpha, \frac{\alpha(2k+2)}{2\sigma_p^2})$ -RDP. So, by standard composition of RDP, we have the total privacy cost at any order $\alpha > 1$ to be: $(\alpha, \alpha \cdot (\frac{kT}{2\sigma^2} + \frac{k+1}{\sigma_p^2}))$. We can obtain the final (ϵ, δ) -joint differential privacy guarantee by optimizing for α , similarly to Appendix A.

Loss function. We minimize the following loss in practice.

$$f(\hat{U}, \hat{V}) = \|P_\Omega(M - \hat{U}\hat{V}^\top)\|_F^2 + \lambda_0 \|\hat{U}\hat{V}^\top\|_F^2 + \lambda \sum_{i=1}^n \frac{c_i^\nu}{Z} \|\hat{U}_i\|^2 + \lambda \sum_{j=1}^m \frac{\tilde{c}_j^\mu}{Z'} \|\hat{V}_j\|^2, \quad (35)$$

where $\lambda_0, \lambda, \mu,$ and ν are hyper-parameters. The loss function used in the description of Algorithm 1 is a special case of (35) where $\lambda_0 = \mu = \nu = 0$. The additional terms in (35) do not change the essence of the algorithm, but we find that they make a significant difference in practice.

First, the term $\lambda_0 \|\hat{U}\hat{V}^\top\|_F^2$ is often used in problems with implicit feedback, as in (Hu et al., 2008). In such problems, the observed entries are often binary, and minimizing the objective $\|P_\Omega(M - \hat{U}\hat{V}^\top)\|_F^2$ can yield a trivial solution – the matrix of all ones. The addition of the second term penalizes non-zero predictions outside of Ω , leading to better generalization. One of the benchmarks we use is an implicit feedback task, in which the use of the second term is necessary. As described in Section 5.2, this results in an additional term \mathbf{K} in Line 9 of $\mathcal{A}_{\text{item}}$, and care is needed when adding privacy protection to

this term, since it involves a sum over all user embeddings. The key observation is that this term is constant for all items, so we only need to compute a noisy version of \mathbf{K} once and use it for all items, thus limiting the privacy loss it incurs.

Second, we use a weighted ℓ_2 regularization, where the weights are defined as follows. The weight of movie j is \tilde{c}_j^μ / Z' , where \tilde{c} is the vector of approximate counts (computed in Algorithm 3), μ is a non-negative hyper-parameter and Z' is the normalizing constant $Z' = \frac{1}{m} \sum_{j=1}^m \tilde{c}_j^\mu$. When μ is positive, this corresponds to applying heavier regularization to more frequent items, and we found in our experiments that this can significantly help generalization. The weights for the users are defined similarly, with one main difference: instead of using approximate counts \tilde{c} , we use the exact counts c , as this term only affects the solution in $\mathcal{A}_{\text{user}}$, which is a privileged computation as illustrated in Figure 1.

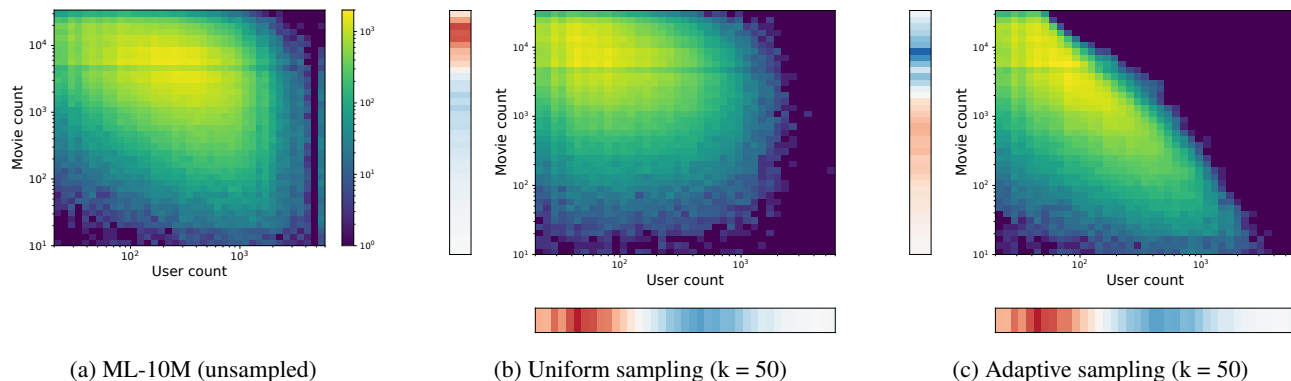


Figure 6. Histogram of user and movie counts in ML-10M, in the original data, and under uniform and adaptive sampling. The color bars in Figures 6b and 6c show the difference in marginal probability compared to the original data in 6a. Red indicates an increase in marginal probability, while blue indicates a decrease. Note that the probability of frequent movies increases under uniform sampling, and decreases under adaptive sampling.

Effect of uniform and adaptive sampling. As observed in Figure 2, the movie count distribution of the MovieLens data set is heavily skewed. We also observed that, perhaps surprisingly, uniformly sampling k items per user tangibly increases the skew. This can be explained by a negative correlation between user counts and movie counts; we computed a correlation coefficient of -0.243 . This is also visible in Figure 6a, which shows the joint histogram of $\{(c_i, c_j), (i, j) \in \Omega\}$, where $c_i = |\Omega_i|$ is the user count (the number of ratings this user produced) and $c_j = |\Omega_j|$ is the movie count (the number of ratings the movie received). The figure illustrates that infrequent users are more likely to rate frequent movies than the average user. By uniformly sampling a constant number of movies per user (Figure 6b), we are, by definition, increasing the probability of infrequent users, hence increasing the probability of frequent movies (due to the negative correlation). This is made clear by the color bar left of Figure 6b, which shows the change in movie count probability with respect to the original data set. This increase in the probability of frequent movies aggravates the skew of the movie distribution, as seen in Figure 2.

Adaptive sampling has the opposite effect: Figure 6c shows that the probability of frequent movies *decreases* under adaptive sampling, while that of infrequent movies increases. This leads to a decrease in bias toward frequent movies, as shown in Figure 2, and results in a significant improvement in the privacy/utility trade-off as discussed in Section 6.3.

D. Additional Details on Experiments

D.1. Details on the Experimental Setup

Table 2 shows the statistics of the MovieLens data sets.

Table 2. Statistics of the experiment data sets.

	ML-10M-top400	ML-10M	ML-20M
n (number of users)	69,692	69,878	136,677
m (number of items)	400	10,677	20,108
$ \Omega $ (number of observations)	4.49M	10M	9.99M

For each data set, we partition the set of observations Ω into $\Omega = \Omega^{\text{train}} \sqcup \Omega^{\text{valid}} \sqcup \Omega^{\text{test}}$. Hyper-parameter tuning is performed on Ω^{valid} , and the final results are reported on Ω^{test} . The pre-processing described in Algorithm 3 is only applied to Ω^{train} .

In the ML-10M benchmark, we follow the setup of (Lee et al., 2013) and use a 80-10-10 split (random uniform over Ω). In the ML-10M-top400 benchmark, we follow the setup of (Jain et al., 2018) and use a 98-1-1 split (random uniform over Ω). In the ML-20M benchmark, we follow the setup of (Liang et al., 2018) and partition the set by *users*, that is, a set of 20K random users are held-out, half of which are used for validation, and the other half for testing. Note that since held-out users are never seen in training, the protocol is to further split each user’s observations Ω_i^{test} (uniformly at random) into $\Omega_i^{\text{test query}} \sqcup \Omega_i^{\text{test target}}$. At test time, the model is allowed access to $\Omega_i^{\text{test query}}$ to compute a user embedding and make a prediction for the user, and $\Omega_i^{\text{test target}}$ is used as the ground truth target. The user embedding \hat{U}_i is computed at test time simply by minimizing the loss in Eq. (35) given the learned movie embeddings \hat{V} , that is,

$$\hat{U}_i = \arg \min_{u \in \mathbb{R}^r} \|P_{\Omega_i^{\text{test query}}}(\mathbf{M}_i - u\hat{V}^\top)\|_F^2 + \lambda_0 \|u\hat{V}^\top\|_F^2 + \lambda \frac{c_i^\nu}{Z'} \|u\|^2.$$

The resulting \hat{U}_i is used to generate predictions for user i . Note that this procedure is consistent with the Joint-DP setting: the computation of \hat{U}_i corresponds to one step of $\mathcal{A}_{\text{user}}$ in Algorithm 1, and is considered privileged (see Figure 1). Besides, since the resulting embedding is not further used for training, it is unnecessary to clip the embedding norm. Avoiding norm clipping at test time could result in better predictions.

Finally the recall for user i is computed as follows. Let $\Omega_i^{\text{prediction}}$ be the top k items that are not in $\Omega_i^{\text{test query}}$. Then $\text{Recall}@k = \frac{|\Omega_i^{\text{prediction}} \cap \Omega_i^{\text{test target}}|}{\min(k, |\Omega_i^{\text{test target}}|)}$.

D.2. Hyper-Parameter Description and Ranges

Table 3 summarizes the complete list of hyper-parameters used in Algorithm 1, Algorithm 3, and in the loss function (35), and specifies the ranges used in our experiments. We make several remarks about hyper-parameters:

Table 3. Hyper-parameter description and ranges.

Symbol	Description	Range
Model and training parameters		
r	rank	[2, 128]
λ	ℓ_2 regularization coefficient	[0.1, 100]
λ_0	coefficient of the global penalty term	[0.1, 5]
μ	item regularization exponent	{0, 0.5, 1}
ν	user regularization exponent	{0, 0.5, 1}
T	number of steps	[1, 5]
Privacy parameters		
Γ_u	row clipping parameter	1
Γ_M	entry clipping parameter	{1, 5}
k	maximum number of ratings per user	[20, 150]
σ	noise standard deviation	see remark below
Pre-processing parameters		
β	fraction of items to train on	[0, 1]
σ_p	standard deviation of pre-processing noise	[10, 200]

- In the non-private baselines, only the model and training parameters are tuned.
- Pre-processing (Algorithm 3) is not used in synthetic experiments. Indeed, these heuristics are designed to deal with the non-uniform distribution of observations in practice. In synthetic experiments, the distribution is uniform by design.
- In the MovieLens experiments, the maximum value in M is known by definition of the task: In ML-10M, entries represent ratings in the interval [0.5, 5], and in ML-20M the entries are binary. Thus, we simply set Γ_M to this value without tuning.

Private Alternating Least Squares

- We find that carefully tuning the model parameters, including the regularization coefficients λ, λ_0 and the exponents μ, ν is important and can have a significant effect.
- For the rating prediction tasks (ML-10M and ML-10M-top400), we find that setting λ_0 to a positive number is detrimental, so we always use 0. For the item recommendation task (ML-20M), using a non-zero λ_0 is important.
- The partitioning of the movies into Frequent and Infrequent is important for the private models, especially at lower values of ε (see Figure 10), but does not help for the non-private baselines.
- To set the standard deviation σ , we use the simple observation that when all hyper-parameters except σ are fixed, ε is a decreasing function of σ that can be computed in closed form. Therefore, in each experiment, we set a target value of ε and do a binary search over σ to select the smallest value that achieves the target ε .
- Finally, note that in Algorithm 1, the parameter σ determines the standard deviation of two noise terms: \mathbf{G} in Line 7 and \mathbf{g} in Line 8. While this is sufficient for the analysis, we find in practice that the model is often more sensitive to \mathbf{g} , thus it can be advantageous to use different scales of noise. We will use the symbols σ_G, σ_g to specify the scales of each term.

The optimal hyper-parameter values for each experiment and each value of ε are given in Table 4. These values are obtained through cross-validation. We do not include the privacy loss of hyper-parameter search because our main objective is to give insights into the choice of hyper-parameters at different privacy budgets. In practice, this can be accounted for, for example by the method in (Liu & Talwar, 2019).

Table 4. Optimal hyper-parameter values for the experiments in Figure 3. The clipping parameter Γ_u is set to 1 in all experiments.

	ML-10M-top400					ML-10M					ML-20M				
	DPALS				ALS	DPALS				ALS	DPALS				ALS
ε	0.8	4	8	16	-	1	5	10	20	-	1	5	10	20	-
r	50	50	50	50	50	32	128	128	128	128	32	32	32	128	128
λ	90	90	80	80	70	120	80	70	60	70	0.5	0.5	0.1	50	30
λ_0	0	0	0	0	0	0	0	0	0	0	2	0.6	0.4	0.4	0.1
μ	0.5	0.5	0.5	0.5	1	0.5	0.5	0.5	0.5	1	-	-	-	-	-
ν	1	1	1	1	1	1	1	1	1	1	-	-	-	-	-
T	2	2	2	2	15	2	2	2	2	15	1	3	3	1	15
k	40	50	50	50	-	50	50	50	50	-	60	60	100	60	-
σ_G	126.9	29.0	11.3	5.86	-	125.9	27.8	15.5	7.5	-	64.0	20.2	14.0	3.5	-
σ_g	63.4	14.5	11.3	5.86	-	63.0	13.9	7.7	3.8	-	64.0	20.2	14.0	3.5	-
β	1	1	1	1	-	0.05	0.4	0.5	0.6	-	0.05	0.1	0.05	0.05	-
σ_p	200	200	20	20	-	100	20	10	10	-	100	100	100	100	-

D.3. Standard Deviation

Finally, Table 5 reports the standard deviation of the DPALS metrics in Figure 3. For each data point, we repeat the experiment 20 times, using the same set of hyper-parameters selected on the validation set, and report the mean and standard deviation of the metric measured on the test set. In all cases, the standard deviation is less than 0.5% of the mean.

Table 5. Mean and standard deviation of the DPALS metrics in Figure 3.

	ML-10M-top400 (test RMSE)				ML-10M (test RMSE)				ML-20M (test Recall@20)			
	0.8	4	8	16	1	5	10	20	1	5	10	20
mean	0.8855	0.8321	0.8201	0.8147	0.9398	0.8725	0.8530	0.8373	0.3120	0.3330	0.3368	0.3444
stddev	0.0025	0.0009	0.0011	0.0008	0.0009	0.0006	0.0004	0.0005	0.0016	0.0010	0.0012	0.0013

D.4. Additional Experiments

Convergence plots for DPALS and DPFW. This experiment illustrates the fact that ALS converges faster than FW, both in its exact and private variants, making it more suitable for training private models. Figure 7 shows the test error (RMSE) against number of iterations, on the synthetic data set with $n = 20K$ users. We use the vanilla version of DPALS without the heuristics introduced in Section 5. The hyper-parameters of both methods are tuned on the validation set.

For the non-private baselines, ALS converges significantly faster than FW. For example, the error of ALS after 2 iterations is lower than the error of FW after 40 iterations. For the private models, we compare the two methods with the same sampling rate ($k = 150$) and same noise level ($\sigma = 10$ in Figure 7a and $\sigma = 20$ in Figure 7b), and tune other parameters. Since the sampling rate and noise level are fixed, the ϵ level is directly determined by the number of steps, and the vertical lines show different levels of ϵ . We can make the following observations. For both methods, in the presence of noise, the error decreases for a few iterations at a rate similar to their exact variants, then plateaus at a fixed error. The fixed error for DPALS is an order of magnitude lower than DPFW. Furthermore, the error reached by DPALS in 2 iterations is lower than the error of DPFW after 40 iterations. The faster convergence of DPALS, even in the presence of noise, directly translates to a better privacy/utility trade-off as demonstrated in Section 6.

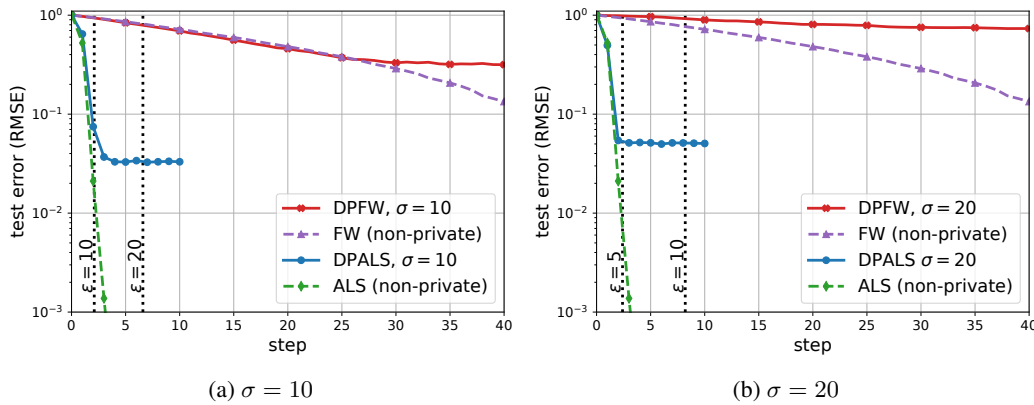


Figure 7. RMSE against steps on the synthetic data set with $n = 20K$. Dashed lines correspond to the non-private baselines (without noise) and solid lines correspond to the private methods with a fixed noise level (left: $\sigma = 10$, right: $\sigma = 20$).

Varying the number of users. This experiment further illustrates the effect of increasing the number of users. We train the DPALS on a subset of the ML-10M-top400 data set, obtained by randomly sampling a subset of n users. Figure 8 shows the results for different values of n , and confirms that increasing the number of users (while keeping the number of movies constant) improves the privacy/utility trade-off. The figure also compares to the DPFW baseline trained on the full data ($n = 69692$). Note that DPALS significantly outperforms DPFW even when trained on a small fraction of the users ($n = 16000$, or 26.4% of the total users).




Figure 8. DPALS on ML-10M-top400 with a varying number of users, n .

Effect of the rank. This experiment explores the effect of the rank on the privacy/utility tradeoff. Figure 9 shows the trade-offs for models of different ranks r on ML-10M and ML-20M. We observe that for non-private ALS, models of higher rank consistently achieve better performance in the range of ranks that we have tried. This is not always the case for the private models. For the ML-10M task, the higher rank model ($r = 128$) performs well for larger values of ϵ , but not for $\epsilon = 1$. On the ML-20M task, the private model with $r = 128$ gives the best recall for $\epsilon \approx 20$ while $r = 32$ performs the best for smaller ϵ . Therefore, unlike in the non-private ALS algorithm where a higher rank is often more desirable given enough computational and storage resources, when training a private model, one needs to carefully choose the rank to balance model capacity and utility degradation due to privacy.




Figure 9. Privacy/utility trade-off for models of different ranks. For lower values of ϵ , a lower rank achieves a better privacy/utility.

Training on Frequent movies. This experiment illustrates the effect of partitioning the set movies into (Frequent \sqcup Infrequent) and training only on Frequent movies. Figure 10 shows the test RMSE vs movie fraction, at different levels of ϵ . The rank of the model is fixed to $r = 32$, the sample size is fixed to $k = 50$, and other hyper-parameters are re-tuned. The results show that as ϵ decreases, the optimal fraction of movies decreases. In particular, for $\epsilon = 1$, the optimal fraction is 5%; note however that this still corresponds to more than 50% of the ratings, as shown on the right sub-figure.




Figure 10. RMSE vs movie fraction, for a rank 32 model on ML-10M, at different privacy levels ϵ . Both figures show the same data, but with a different x axis. The movie fraction (left figure) is defined as $|\text{Frequent}|/m$. The data fraction (right figure) is defined as $|\{(i, j) \in \Omega : j \in \text{Frequent}\}|/|\Omega|$. The right figure emphasizes the long-tail distribution of movie counts – a small fraction of Frequent movies corresponds to a large fraction of data.

Figure 11 shows a similar result for ML-20M. The optimal movie fraction in this example is between 5% and 10% depending on the rank.




Figure 11. Recall@20 vs movie fraction on ML-20M, for $\epsilon = 5$.

Effect of the regularization exponents. This experiment illustrates the effect of the regularization exponents (ν, μ) in the loss function (35). We vary (ν, μ) for a rank 128 model with $\epsilon = 10$ on ML-10M (and re-tune other parameters). The results are reported in Figure 12. This example indicates that a careful tuning of the ℓ_2 regularization can have a significant impact on utility, and can also make the private models more robust to noise: Notice that with the optimal setting of (ν, μ) the model can be trained on a much larger fraction of movies, with only a slight degradation in utility.




Figure 12. RMSE vs movie fraction on ML-10M, for $\epsilon = 10$ and $r = 128$, and for different values of regularization exponents (ν, μ) .