Appendix: Active Positive Semidefinite Matrix Completion: Algorithms, Theory and Applications

immediate

March 1, 2017

Abstract

We provide proofs that were skipped in the main paper. We also provide some additional experimental results and related work concerning multi-armed bandits that was skipped in the main paper.

1 Preliminaries

We shall repeat a proposition that was stated in the main paper for the sake of completeness.

Proposition 1.1. *Let $\boldsymbol{L}$ be any SPSD matrix of size K. Given a subset $\mathcal{C} \subset \{1, 2, \dots, K\}$, the columns of the matrix $\boldsymbol{L}$ indexed by the set $\mathcal{C}$ are independent iff the principal submatrix $\boldsymbol{L}_{\mathcal{C},\mathcal{C}}$ is non-degenerate, equivalently iff, $\lambda_{\min}(\boldsymbol{L}_{\mathcal{C},\mathcal{C}}) > 0$.*

We would also need the classical matrix Bernstein inequality, which we borrow from the work of Joel Tropp [Tropp, 2015].

Theorem 1.2. *Let $\boldsymbol{S}_1, \dots, \boldsymbol{S}_n$ be independent, centered random matrices with dimension $d_1 \times d_2$ and assume that each one is uniformly bounded*

$$\mathbb{E}\boldsymbol{S}_k = 0, \|\boldsymbol{S}_k\| \leq L \text{ for each } k = 1, \dots, n.$$

Introduce the sum $\boldsymbol{Z} = \sum_{k=1}^{n} \boldsymbol{S}_k$, and let $\nu(\boldsymbol{Z})$ denote the matrix variance statistic of the sum:

$$\nu(\boldsymbol{Z}) = \max\left\{\left\|\mathbb{E}\boldsymbol{Z}\boldsymbol{Z}^\top\right\|, \left\|\mathbb{E}\boldsymbol{Z}^\top\boldsymbol{Z}\right\|\right\} \tag{1}$$

$$= \max\left\{\left\|\sum_{k=1}^{n}\mathbb{E}\boldsymbol{S}_k\boldsymbol{S}_k^\top\right\|, \left\|\sum_{k=1}^{n}\mathbb{E}\boldsymbol{S}_k^\top\boldsymbol{S}_k\right\|\right\} \tag{2}$$

Then,

$$\mathbb{P}(\|\boldsymbol{Z}\| \geq t) \leq (d_1 + d_2)\exp\left(-\frac{t^2/2}{\nu(\boldsymbol{Z}) + \frac{Lt}{3}}\right)$$

2 Sample complexity of MCANS algorithm: Proof of Theorem 3.2 in the main paper

Theorem 2.1. *If $\boldsymbol{L} \in \mathbb{R}^{K\times K}$ is an SPSD matrix of rank r, then the matrix $\hat{\boldsymbol{L}}$ output by the MCANS algorithm satisfies $\hat{\boldsymbol{L}} = \boldsymbol{L}$. Moreover, the number of oracle calls made by MCANS is at most $K(r+1)$. The sampling algorithm requires: $K + (K-1) + (K-2) + \ldots + (K-(r-1)) + (K-r) \leq (r+1)K$ samples from the matrix $\boldsymbol{L}$.*

Proof. MCANS checks one column at a time starting from the second column, and uses the test in line 5 to determine if the current column is independent of the previous columns. The validity of this test is guaranteed by proposition (1.1). Each such test needs just one additional sample corresponding to the index (c, c). If a column c is found to be independent of the columns $1, 2, \ldots, c-1$ then rest of the entries in column c are queried. Notice, that by now we have already queried all the columns and rows of matrix $\boldsymbol{L}$ indexed by the set $\mathcal{C}$, and also queried the element (c, c) in line 4. Hence we need to query only $K - |\mathcal{C}| - 1$ more entries in column c in order to have all the entries of column c. Combined with the fact that we query only r columns completely and in the worst case all the diagonal entries might be queried, we get the total query complexity to be $(K-1) + (K-2) + \ldots (K-r) + K \leq K(r+1)$. □

3 Proof of Lemma 4.1 in the main paper

We begin by stating the lemma.

Lemma. *Let $\hat{\boldsymbol{P}}$ be a $p \times p$ random matrix that is constructed as follows. For each index (i, j) independent of other indices, set $\hat{\boldsymbol{P}}_{i,j} = \frac{H_{i,j}}{n_{i,j}}$, where $H_{i,j}$ is a random variable drawn from the distribution $\mathrm{Binomial}(n_{i,j}, p_{i,j})$. Let $Z = \hat{\boldsymbol{P}} - \boldsymbol{P}$. Then,*

$$||Z||_2 \leq \frac{2\log(2p/\delta)}{3 \min\limits_{i,j} n_{i,j}} + \sqrt{\frac{\log(2p/\delta)}{2} \sum_{i,j} \frac{1}{n_{i,j}}}. \tag{3}$$

Furthermore, if we denote by Δ the R.H.S. in Equation (3), then $|\sigma_{\min}(\hat{\boldsymbol{P}}) - \sigma_{\min}(\boldsymbol{P})| \leq \Delta$.

Proof. Define, $\boldsymbol{S}^t_{i,j} = \frac{1}{n_{i,j}}(X^t_{i,j} - p_{i,j})\boldsymbol{E}_{i,j}$, where $\boldsymbol{E}_{i,j}$ is a $p \times p$ matrix with a 1 in the $(i, j)^{\text{th}}$ entry and 0 everywhere else, and $X^t_{i,j}$ is a random variable sampled from the distribution $\mathrm{Bern}(p_{i,j})$. If $X^t_{i,j}$ are independent for all t, i, j, then it is easy to see that $Z = \sum_{i,j} \frac{1}{n_{i,j}} \sum_{t=1}^{n_{i,j}} \boldsymbol{S}^t_{i,j}$. Hence $\boldsymbol{S}$ is a sum of independent random matrices and this allows to apply matrix Bernstein type inequalities. In order to apply the matrix Bernstein inequality, we would need upper bound on maximum spectral norm of the summands, and an upper bound on the variance of $\boldsymbol{Z}$. We next bound these two quantities as follows,

$$||\boldsymbol{S}^t_{i,j}||_2 = ||\frac{1}{n_{i,j}}(X^t_{i,j} - p_{i,j})\boldsymbol{E}_{i,j}||_2 = \frac{1}{n_{i,j}}|X^t_{i,j} - p_{i,j}| \leq \frac{1}{n_{i,j}}. \tag{4}$$

To bound the variance of $\boldsymbol{Z}$ we proceed as follows

$$\nu(Z) = ||\sum_{i,j}\sum_{t=1}^{n_{i,j}} \mathbb{E}(\boldsymbol{S}^t_{i,j})^\top \boldsymbol{S}^t_{i,j}|| \bigwedge ||\sum_{i,j}\sum_{t=1}^{n_{i,j}} \mathbb{E}\boldsymbol{S}^t_{i,j}(\boldsymbol{S}^t_{i,j})^\top|| \tag{5}$$

Via elementary algebra and using the fact that $\mathrm{Var}(X^t_{i,j}) = p_{i,j}(1 - p_{i,j})$It is easy to see that,

$$\mathbb{E}(\boldsymbol{S}^t_{i,j})^\top \boldsymbol{S}^t_{i,j} = \frac{1}{n^2_{i,j}} \mathbb{E}(X^t_{i,j} - p_{i,j})^2 (\boldsymbol{E}^t_{i,j})^\top \boldsymbol{E}^t_{i,j} \tag{6}$$

$$= \frac{1}{4n^2_{i,j}} \boldsymbol{E}_{i,i}. \tag{7}$$

Using similar calculations we get $\mathbb{E}\boldsymbol{S}^t_{i,j}(\boldsymbol{S}^t_{i,j})^\top = \frac{1}{4n^2_{i,j}}\boldsymbol{E}_{j,j}$. Hence, $\nu(Z) = \sum\limits_{i,j} \sum_{t=1}^{n_{i,j}} \frac{1}{4n^2_{i,j}} = \sum_{i,j} \frac{1}{4n_{i,j}}$. Applying matrix Bernstein, we get with probability at least $1 - \delta$

$$||Z||_2 \leq \frac{2\log(2p/\delta)}{3 \min\limits_{i,j} n_{i,j}} + \sqrt{\frac{\log(2p/\delta)}{2} \sum_{i,j} \frac{1}{n_{i,j}}}. \tag{8}$$

The second part of the result follows immediately from Weyl's inequality which says that $|\sigma_{\min}(\hat{\boldsymbol{P}}) - \sigma_{\min}(\boldsymbol{P})| \leq ||\hat{\boldsymbol{P}} - \boldsymbol{P}|| = ||Z||$. □

4 Sample complexity of successive elimination algorithm: Proof of Lemma 4.2 in the main paper

Lemma. *The successive elimination algorithm shown in Figure (6.2) on m square matrices of size $\boldsymbol{A}_1, \ldots, \boldsymbol{A}_m$ each of size $p \times p$ outputs an index $i_\star$ such that, with probability at least $1-\delta$, the matrix $\boldsymbol{A}_{i_\star}$ has the largest smallest singular value among all the input matrices. The total number of queries to the stochastic oracle are*

$$\sum_{k=2}^{m} O\left(\frac{p^3 \log(2p\pi^2 m^2/3\Delta_k^2\delta)}{\Delta_k^2}\right) + O\left(p^4 \max_k \left(\frac{\log(2p\pi^2 m^2/3\Delta_k^2\delta)}{\Delta_k^2}\right)\right) \tag{9}$$

where $\Delta_{k,p} := \max_{j=1,\ldots,m} \sigma_{\min}(\boldsymbol{A}_j) - \sigma_{\min}(\boldsymbol{A}_k)$

Proof. Suppose matrix $\boldsymbol{A}_1$ has the largest smallest singular value. From lemma (3), we know that with probability at least $1-\delta_t$, $|\sigma_{\min}(\widehat{\boldsymbol{A}}_k) - \sigma_{\min}(\boldsymbol{A}_k)| \le \frac{2\log(2p/\delta_t)}{3 \min_{i,j} n_{i,j}(\boldsymbol{A})} + \sqrt{\frac{\log(2p/\delta_t)}{2} \sum_{i,j} \frac{1}{n_{i,j}(\boldsymbol{A})}}$. Hence, by union bound the probability that the matrix $\boldsymbol{A}_1$ is eliminated in one of the rounds is at most $\sum_t \sum_{k=1}^m \delta_t \le \sum_{t=1}^{\max} \sum_{k=1}^m \frac{6\delta}{\pi^2 m t^2} = \delta$. This proves that the successive elimination step identifies the matrix with the largest smallest singular value.

An arm k is eliminated in round t if $\alpha_{t,1} + \alpha_{t,k} \le \hat{\sigma}_t^{\max} - \sigma_{\min}(\widehat{\boldsymbol{A}}_k)$. By definition,

$$\Delta_{k,p} - (\alpha_{t,1} + \alpha_{t,k}) = (\sigma_{\min}(\boldsymbol{A}_1) - \alpha_{t,1}) - (\sigma_{\min}(\boldsymbol{A}_k) + \alpha_{t,k}) \ge \sigma_{\min}(\hat{\boldsymbol{A}}_1) - \sigma_{\min}(\boldsymbol{A}_k) \ge \alpha_{t,1} + \alpha_{t,k} \tag{10}$$

That is if $\alpha_{t,1} + \alpha_{t,k} \le \frac{\Delta_{k,p}}{2}$, then arm k is eliminated in round t. By construction, since in round t each element in each of the surviving set of matrices has been queried at least t times, we can say that $\alpha_{t,j} \le \frac{2\log(2p/\delta_t)}{3t} + \sqrt{\frac{p^2 \log(2p/\delta_t)}{2t}}$ for any index j corresponding to the set of surviving arms. Hence arm k gets eliminated after

$$t_k = O\left(\frac{p^2 \log(2p\pi^2 m^2/3\Delta_{k,p}^2\delta)}{\Delta_{k,p}^2}\right) \tag{11}$$

In each round t the number of queries made are $O(p)$ for each of the m matrices corresponding to the row and column which is different among them, and $O(p^2)$ corresponding to the left $p-1 \times p-1$ submatrix that is common to all of the matrices $A_1, \ldots, A_m$. Hence, the total number of queries to the stochastic oracle is

$$p\sum_{k=2}^{m} t_k + p^2 \max_k t_k = \sum_{k=2}^{m} O\left(\frac{p^3 \log(2p\pi^2 m^2/3\Delta_{k,p}^2\delta)}{\Delta_{k,p}^2}\right) + O\left(p^4 \max_k \left(\frac{\log(2p\pi^2 m^2/3\Delta_{k,p}^2\delta)}{\Delta_{k,p}^2}\right)\right) \quad \square$$

5 Proof of Nystrom method

In this supplementary material we provide a proof of Nystrom extension in $\max$ norm when we use a stochastic oracle to obtain estimators $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ of the matrices $\boldsymbol{C}, \boldsymbol{W}$. The question that we are interested in is how good is the estimate of the Nystrom extension obtained using matrices $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ w.r.t. the Nystrom extension obtained using matrices $\boldsymbol{C}, \boldsymbol{W}$. This is answered in the theorem below.

Theorem 5.1. *Suppose the matrix $\boldsymbol{W}$ is an invertible $r \times r$ matrix. Suppose, by multiple calls to a stochastic oracle we construct estimators $\widehat{\boldsymbol{C}}, \widehat{\boldsymbol{W}}$ of $\boldsymbol{C}, \boldsymbol{W}$. Now, consider the matrix $\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top$ as an estimate $\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top$. Given any $\delta \in (0,1)$, with probability atleast $1-\delta$,*

$$\left\| \boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top \right\|_{\max} \le \epsilon$$

after making M number of oracle calls to a stochastic oracle, where

$$M \ge 100 C_1(W,C)\log(2Kr/\delta)\max\left(\frac{Kr^{7/2}}{\epsilon}, \frac{Kr^3}{\epsilon^2}\right) + 200 C_2(W,C)\log(2r/\delta)\max\left(\frac{r^5}{\epsilon}, \frac{r^7}{\epsilon^2}\right)$$

where $C_1(\boldsymbol{W}, \boldsymbol{C})$ and $C_2(\boldsymbol{W}, \boldsymbol{C})$ are given by the following equations

$$C_1(\boldsymbol{W}, \boldsymbol{C}) = \max\left(\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\|_{\max}, \|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\|_{\max}^2, \|\boldsymbol{W}^{-1}\|_{\max}, \|\boldsymbol{C}\boldsymbol{W}^{-1}\|_1^2, \|\boldsymbol{W}^{-1}\|_2 \|\boldsymbol{W}^{-1}\|_{\max}\right)$$
$$C_2(\boldsymbol{W}, \boldsymbol{C}) = \max\left(\|\boldsymbol{W}^{-1}\|_2^2 \|\boldsymbol{W}^{-1}\|_{\max}^2, \|\boldsymbol{W}^{-1}\|_2 \|\boldsymbol{W}^{-1}\|_{\max}, \|\boldsymbol{W}^{-1}\|_2, \|\boldsymbol{W}^{-1}\|_2^2\right)$$

Our proof proceeds by a series of lemmas, which we state next.

Lemma 5.2.

$$\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \le \left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C} - \widehat{\boldsymbol{C}})^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max}$$

Proof.

$$\begin{aligned}
\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} &= \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top + \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&\le \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&= \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \\
&\quad \left\|\widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top + \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&\le \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}W^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \\
&\quad \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\boldsymbol{C}^\top - \widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}\widehat{\boldsymbol{C}}^\top\right\|_{\max} \\
&= \left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C} - \widehat{\boldsymbol{C}})^\top\right\|_{\max} + \left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1} - \widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max}
\end{aligned}$$

□

In the following lemmas we shall bound the three terms that appear in the R.H.S of the bound of Lemma (5.2).

Lemma 5.3.

$$\left\|(C - \hat{C})W^{-1}C^\top\right\|_{\max} \le \frac{2||\boldsymbol{W}^{-1}\boldsymbol{C}^\top||_{\max}}{3m}\log(2Kr/\delta) + \sqrt{\frac{r\,||\boldsymbol{W}^{-1}\boldsymbol{C}^\top||_{\max}^2 \log(2Kr/\delta)}{2m}} \tag{12}$$

Proof. Let $\boldsymbol{M} = \boldsymbol{W}^{-1}\boldsymbol{C}^\top$, then $\left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} = \left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})\boldsymbol{M}\right\|_{\max}$. By the definition of max norm we have

$$\left\|(\boldsymbol{C} - \widehat{\boldsymbol{C}})M\right\|_{\max} = \max_{i,j}\left|\sum_{p=1}^{l}(\boldsymbol{C} - \widehat{\boldsymbol{C}})_{i,p}\boldsymbol{M}_{p,j}\right|$$

Fix a pair of indices (i, j), and consider the expression $\left|\sum_{p=1}^{l}(\boldsymbol{C} - \widehat{\boldsymbol{C}})_{i,p}\boldsymbol{M}_{p,j}\right|$

Define $r_{i,p} = (\boldsymbol{C} - \widehat{\boldsymbol{C}})_{i,p}$. By definition of $r_{i,p}$ we can write $r_{i,p} = \frac{1}{m}\sum_{t=1}^{m} r_{i,p}^t$, where $r_{i,p}^t$ are a set of independent random variables with mean 0 and variance at most $1/4$. This decomposition combined with scalar Bernstein inequality gives that with probability at least $1 - \delta$

$$\begin{aligned}
\left|\sum_{p=1}^{l}(\widehat{\boldsymbol{C}} - \widehat{\boldsymbol{C}})_{i,p}\boldsymbol{M}_{p,j}\right| &= \left|\sum_{p=1}^{l} r_{i,p}\boldsymbol{M}_{p,j}\right| \\
&= \left|\sum_{p=1}^{l}\sum_{t=1}^{m}\frac{1}{m}r_{i,p}^t\boldsymbol{M}_{p,j}\right| \\
&\le \frac{2||M||_{\max}}{3m}\log(2/\delta) + \sqrt{\frac{r\,||M||_{\max}^2\log(2/\delta)}{2m}}
\end{aligned}$$

Applying a union bound over all possible Kr choices of index pairs (i, j), we get the desired result. $\square$

Before we establish bounds on the remaining two terms in the RHS of Lemma (5.2) we state and prove a simple proposition that will be used at many places in the rest of the proof.

Proposition 5.4. *For any two real matrices $M_1 \in \mathbb{R}^{n_1 \times n_2}, M_2 \in \mathbb{R}^{n_2 \times n_3}$ the following set of inequalities are true:*

1. $\|\boldsymbol{M}_1 \boldsymbol{M}_2\|_{\max} \le \|\boldsymbol{M}_1\|_{\max} \|\boldsymbol{M}_2\|_1$
2. $\|\boldsymbol{M}_1 \boldsymbol{M}_2\|_{\max} \le \left\|\boldsymbol{M}_1^\top\right\|_1 \|\boldsymbol{M}_2\|_{\max}$
3. $\|\boldsymbol{M}_1 \boldsymbol{M}_2\|_{\max} \le \|\boldsymbol{M}_1\|_2 \|\boldsymbol{M}_2\|_{\max}$
4. $\|\boldsymbol{M}_1 \boldsymbol{M}_2\|_{\max} \le \|\boldsymbol{M}_2\|_2 \|\boldsymbol{M}_1\|_{\max}$

where, the $\|\cdot\|_p$ is the induced p norm.

Proof. Let $\boldsymbol{e}_i$ denote the i^{th} canonical basis vectors in $\mathbb{R}^K$. We have,

$$\begin{aligned}
\|\boldsymbol{M}_1 \boldsymbol{M}_2\|_{\max} &= \max_{i,j} \left|e_i^\top \boldsymbol{M}_1 \boldsymbol{M}_2 e_j\right| \\
&\le \max_{i,j} \left\|e_i^\top \boldsymbol{M}_1\right\|_{\max} \|\boldsymbol{M}_2 e_j\|_1 \\
&= \max_i \left\|e_i^\top \boldsymbol{M}_1\right\|_{\max} \max_i \|\boldsymbol{M}_2 e_j\|_1 \\
&= \|\boldsymbol{M}_1\|_{\max} \|\boldsymbol{M}_2\|_1 .
\end{aligned}$$

To obtain the first inequality above we used Holder's inequality and the last equality follows from the definition of $||\cdot||_1$ norm. To get the second inequality, we use the observations that $\|\boldsymbol{M}_1 \boldsymbol{M}_2\|_{\max} = \left\|\boldsymbol{M}_2^\top \boldsymbol{M}_1^\top\right\|_{\max}$. Now applying the first inequality to this expression we get the desired result. Similar techniques yield the other two inequalities. $\square$

Lemma 5.5. *With probability at least $1-\delta$, we have*

$$\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} \le \frac{r^2}{2m}\left(\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} + \left\|\boldsymbol{W}^{-1}\right\|_{\max}\right)\log(2Kr/\delta) + \\ r^2\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max}\sqrt{\frac{\log(2Kr/\delta)}{2m}} + r\left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1\sqrt{\frac{\log(2Kr/\delta)}{2m}}$$

Proof.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} &\le \left\|(\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}+\boldsymbol{C}\boldsymbol{W}^{-1})(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} \\
&\overset{(a)}{\le} \left\|(\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1})(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} + \left\|\boldsymbol{C}\boldsymbol{W}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} \\
&\overset{(b)}{\le} \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max}\left\|(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_1 + \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max}\left\|(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_1 \qquad (13)
\end{aligned}$$

To obtain inequality (a) we used triangle inequality for matrix norms, and to obtain inequality (b) we used Proposition (5.4). We next upper bound the first term in the R.H.S. of Equation (13).

We bound the term $\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max}$ next.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max} &\le \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}+\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max} \\
&\le \left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \left\|\boldsymbol{C}\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_{\max} \\
&= \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \left\|\boldsymbol{C}(\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1})\right\|_{\max} \\
&\overset{(a)}{\le} \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\left\|\widehat{\boldsymbol{W}}^{-1}\right\|_{\max} + \left\|\boldsymbol{C}^\top\right\|_1\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} \qquad (14)
\end{aligned}$$

We used Proposition (5.4) to obtain inequality (a). Combining Equations (13) and (14) we get,

$$\begin{aligned}\left\|\widehat{\boldsymbol{C}}\widehat{\boldsymbol{W}}^{-1}(\boldsymbol{C}-\widehat{\boldsymbol{C}})^\top\right\|_{\max} &\leq \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\left(\left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\left\|\widehat{\boldsymbol{W}}^{-1}\right\|_{\max}+\|\boldsymbol{C}^\top\|_1\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max}+\|\boldsymbol{C}\boldsymbol{W}^{-1}\|_{\max}\right)\\ &= \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1^2\left\|\widehat{\boldsymbol{W}}^{-1}\right\|_{\max}+\left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\|\boldsymbol{C}^\top\|_1\left\|\widehat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max}+\\ &\quad \left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1\|\boldsymbol{C}\boldsymbol{W}^{-1}\|_{\max} \qquad (15)\end{aligned}$$

Since all the entries of the matrix $\boldsymbol{C}$ are probabilities we have $\|\boldsymbol{C}\|_{\max} \leq 1$ and $\left\|\boldsymbol{C}^\top\right\|_1 \leq r$. Moreover, since each entry of the matrix $\hat{\boldsymbol{C}}-\boldsymbol{C}$ is the average of m independent random variables with mean 0, and each bounded between $[-1,1]$, by Hoeffding's inequality and union bound, we get that with probability at least $1-\delta$

$$\left\|(\widehat{\boldsymbol{C}}-\boldsymbol{C})^\top\right\|_1 \leq r\sqrt{\frac{\log(2Kr/\delta)}{2m}} \qquad (16)$$

□

The next proposition takes the first steps towards obtaining an upper bound on $\left\|\hat{C}(W^{-1}-\hat{W}^{-1})C^\top\right\|_{\max}$

Proposition 5.6.

$$\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} \leq \min\left\{r^2\left\|\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1}\right\|_{\max}, r\left\|\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1}\right\|_1\right\}$$

Proof.

$$\begin{aligned}\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} &\overset{(a)}{\leq} \left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\right\|_{\max}\|\boldsymbol{C}^\top\|_1\\ &\overset{(b)}{\leq} r\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\right\|_{\max}\\ &\overset{(c)}{\leq} \min\left\{r^2\left\|\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1}\right\|_{\max}, r\left\|\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1}\right\|_1\right\} \qquad (17)\end{aligned}$$

In the above bunch of inequalities (a) and (c) we used Proposition (5.4) and to obtain inequality (b) we used the fact that $||\boldsymbol{C}^\top||_{\max} \leq r$. □

Hence, we need to bound $\left\|\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1}\right\|_{\max}$ and $\left\|\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1}\right\|_1$.

Let us define $\hat{\boldsymbol{W}} = \boldsymbol{W} + \boldsymbol{E}_W$ where $\boldsymbol{E}_W$ is the error-matrix and $\hat{\boldsymbol{W}}$ is the sample average of m independent samples of a random matrix where $\mathbb{E}\hat{\boldsymbol{W}}_k(i,j) = \boldsymbol{W}(i,j)$.

Lemma 5.7. *Let us define $\hat{\boldsymbol{W}} - \boldsymbol{W} = \boldsymbol{E}_W$. Suppose, $\left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2 \leq \frac{1}{2}$, then*

$$\left\|\hat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} \leq 2\left\|\boldsymbol{W}^{-1}\right\|_2\|\boldsymbol{E}_W\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}$$

.

Proof. Since $\left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2 < 1$, we can apply the Taylor series expansion:

$$(\boldsymbol{W}+\boldsymbol{E}_W)^{-1} = \boldsymbol{W}^{-1} - \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \cdots$$

Therefore:

$$\begin{aligned}\left\|\hat{\boldsymbol{W}}^{-1}-\boldsymbol{W}^{-1}\right\|_{\max} &= \|\boldsymbol{W}^{-1} - \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1} + \cdots - \boldsymbol{W}^{-1}\|_{\max}\\ &\overset{(a)}{\leq} \|\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\|_{\max} + \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\boldsymbol{E}_W\boldsymbol{W}^{-1}\right\|_{\max} + \cdots\\ &\overset{(b)}{\leq} \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2\|W^{-1}\|_{\max} + \left\|\boldsymbol{W}^{-1}\boldsymbol{E}_W\right\|_2^2\|W^{-1}\|_{\max} + \dots\\ &\overset{(c)}{\leq} 2\left\|\boldsymbol{W}^{-1}\right\|_2\|\boldsymbol{E}_W\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}\end{aligned}$$

To obtain the last inequality we used the hypothesis of the lemma, and to obtain inequality (a) we used the triangle inequality for norms, and to obtain inequality (b) we used proposition (5.4). Inequlaity (c) follows from the triangle inequality. □

Thanks to Lemma (5.7) and proposition (5.6) we know that $\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} \leq r^2\epsilon$. We now need to guarantee that the hypothesis of lemma (5.7) applies. The next lemma helps in doing that.

Lemma 5.8. *With probability at least $1-\delta$ we have*

$$\|\boldsymbol{E}_W\| = \left\|\hat{\boldsymbol{W}}-\boldsymbol{W}\right\| \leq \frac{2r}{3m}\log(2r/\delta)+\sqrt{\frac{r\log(2r/\delta)}{2m}} \tag{18}$$

Proof. The proof is via matrix Bernstein inequality. By the definition of $\hat{\boldsymbol{W}}$, we know that $\hat{\boldsymbol{W}}-\boldsymbol{W} = \frac{1}{m}\sum(\boldsymbol{W}_i-\boldsymbol{W})$, where $\hat{\boldsymbol{W}}$ is $0-1$ random matrix where the $(i,j)^{\text{th}}$ entry of the matrix $\hat{\boldsymbol{W}}$ is a single Bernoulli sample sampled from $\text{Bern}(\boldsymbol{W}_{i,j})$. For notational convenience denote $Z_i := \frac{1}{m}\hat{\boldsymbol{W}}_i - \boldsymbol{W}$ This makes $\hat{\boldsymbol{W}}-\boldsymbol{W} = \frac{1}{m}\sum \boldsymbol{W}_i - \boldsymbol{W}$ an average of m independent random matrices each of whose entry is a 0 mean random variable with variance at most $1/4$, with each entry being in $[-1,1]$. In order to apply the matrix Bernstein inequality we need to upper bound ν, L (see Theorem (1.2)), which we do next.

$$\left\|\frac{1}{m}(\hat{\boldsymbol{W}}_i-\boldsymbol{W})\right\|_2 \leq \frac{1}{m}\sqrt{r^2} = \frac{r}{m}. \tag{19}$$

In the above inequality we used the fact that each entry of $(\hat{\boldsymbol{W}}_i-\boldsymbol{W})$ is between $[-1,1]$ and hence the spectral norm of this matrix is at most $\sqrt{r^2}$. We next bound the parameter ν.

$$\nu = \frac{1}{m^2}\max\left\{\left\|\sum_i \mathbb{E}\boldsymbol{Z}_i\boldsymbol{Z}_i^\top\right\|, \left\|\sum_i \mathbb{E}\boldsymbol{Z}_i^\top\boldsymbol{Z}_i\right\|\right\} \tag{20}$$

It is not hard to see that the matrix $\mathbb{E}\boldsymbol{Z}_i\boldsymbol{Z}_i^\top$ is a diagonal matrix, where each diagonal entry is at most $\frac{l}{4}$. The same holds true for $\mathbb{E}\boldsymbol{Z}_i\boldsymbol{Z}_i^\top$. Putting this back in Equation (20) we get $\nu \leq \frac{r}{4m}$. Putting $L=\frac{r}{m}$ and $\nu=\frac{r}{4m}$, we get

$$\left\|\hat{\boldsymbol{W}}-\boldsymbol{W}\right\| \leq \frac{2r}{3m}\log(2r/\delta)+\sqrt{\frac{r\log(2r/\delta)}{2m}} \tag{21}$$

□

We are now ready to establish the following bound

Lemma 5.9. *Assuming that $m \geq m_0 := \frac{4r\|\boldsymbol{W}^{-1}\|}{3} + 2r\log(2r/\delta)\left\|\boldsymbol{W}^{-1}\right\|_2^2$, with probability at least $1-\delta$ we will have*

$$\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} \leq 2r^2\left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}\left(\frac{2r}{3m}\log(2r/\delta)+\sqrt{\frac{r\log(2r/\delta)}{2m}}\right). \tag{22}$$

Proof.

$$\begin{aligned}
\left\|\widehat{\boldsymbol{C}}(\boldsymbol{W}^{-1}-\widehat{\boldsymbol{W}}^{-1})\boldsymbol{C}^\top\right\|_{\max} &\overset{\text{(a)}}{\leq} r^2\left\|\boldsymbol{W}^{-1}-\hat{\boldsymbol{W}}^{-1}\right\|_{\max} \\
&\overset{\text{(b)}}{\leq} 2r^2\left\|\boldsymbol{W}^{-1}\boldsymbol{E}_{\boldsymbol{W}}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max} \\
&\overset{\text{(c)}}{\leq} 2r^2\left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{E}_{\boldsymbol{W}}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max} \\
&\overset{\text{(d)}}{\leq} 2r^2\left\|\boldsymbol{W}^{-1}\right\|_2\left\|\boldsymbol{W}^{-1}\right\|_{\max}\left(\frac{2r}{3m}\log(2r/\delta)+\sqrt{\frac{r\log(2r/\delta)}{2m}}\right)
\end{aligned}$$

□

To obtain inequality (a) above we used proposition (5.6), to obtain inequality (b) we used lemma (5.7), and finally to obtain inequality (c) we used the fact that matrix 2-norms are submultiplicative.

With this we now have bounds on all the necessary quantities. The proof of our theorem essentially requires us to put all these terms together.

6 Proof of Theorem 4.3 in the main paper

Since we need the total error in $\max$ norm to be at most ϵ, we will enforce that each term of our expression be atmost $\frac{\epsilon}{10}$. From lemma (5.2) we know that the maxnorm is the sum of three terms. Let us call the three terms in the R.H.S. of Lemma (5.2) T_1, T_2, T_3 repsectively. We then have that if we have m_1 number of copies of the matrix $\boldsymbol{C}$, where

$$m_1 \geq \frac{20 \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max} \log(2Kr/\delta)}{3\epsilon} \bigwedge \frac{100r \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}^2 \log(2Kr/\delta)}{2\epsilon^2} \tag{23}$$

then $T_1 \leq \epsilon/5$. Next we look at T_3. From lemma (5.9) it is easy to see that we need m_3 independent copies of the matrix $\boldsymbol{W}$ so that $T_3 \leq \epsilon/5$, where m_3 is equal to

$$m_3 \geq \frac{40r^3 \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max} \log(2r/\delta)}{3\epsilon} \bigwedge \frac{400r^5 \left\|\boldsymbol{W}^{-1}\right\|_2^2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}^2 \log(2r/\delta)}{2\epsilon^2} \tag{24}$$

Finally we now look at T_2. Combining lemma (5.5), and lemma (5.7) and (5.8) and after some elementary algebraic calculations we get that we need m_2 independent copies of the matrix $\boldsymbol{C}$ and $\boldsymbol{W}$ to get $T_2 \leq \frac{3\epsilon}{5}$, where m_2 is

$$m_2 \geq 100 \max(\left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}) \log(2Kr/\delta) \left(\frac{r^{5/2}}{\epsilon}, \frac{r^2}{\epsilon^2}\right) \tag{25}$$

The number of calls to stochastic oracle is $r^2(m_0 + m_3) + Kr(m_1 + m_2)$, where m_0 is the number as stated in Lemma (5.9). Using the above derived bounds for $m_0 + m_1, m_2, m_3$ we get

$$Kr(m_1 + m_2) + r^2(m_0 + m_3) \geq 100 \log(2Kr/\delta) C_1(W, C) \max\left(\frac{Kr^{7/2}}{\epsilon}, \frac{Kr^3}{\epsilon^2}\right) + \\ 200 C_2(W, C) \log(2r/\delta) \max\left(\frac{r^5}{\epsilon}, \frac{r^7}{\epsilon^2}\right)$$

where $C_1(\boldsymbol{W}, \boldsymbol{C})$ and $C_2(\boldsymbol{W}, \boldsymbol{C})$ are given by the following equations

$$C_1(\boldsymbol{W}, \boldsymbol{C}) = \max\left(\left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\boldsymbol{C}^\top\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{C}\boldsymbol{W}^{-1}\right\|_1^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}\right)$$
$$C_2(\boldsymbol{W}, \boldsymbol{C}) = \max\left(\left\|\boldsymbol{W}^{-1}\right\|_2^2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}^2, \left\|\boldsymbol{W}^{-1}\right\|_2 \left\|\boldsymbol{W}^{-1}\right\|_{\max}, \left\|\boldsymbol{W}^{-1}\right\|_2, \left\|\boldsymbol{W}^{-1}\right\|_2^2\right)$$

7 Additional experimental results: Comparison with LRMC on Movie Lens datasets

First we present the results on the synthetic dataset. To generate a low-rank matrix, we take a random matrix in $\boldsymbol{L}_1 = [0, 1]^{K \times r}$ and then define $\boldsymbol{L}_2 = \boldsymbol{L}_1 \boldsymbol{L}_1^\top$. Then get $\boldsymbol{L} = \boldsymbol{L}_2 / \max_{i,j}(\boldsymbol{L}_2)_{i,j}$. This matrix $\boldsymbol{L}$ will be $K \times K$ and have rank r.

In Figure 2, you can find the comparison of LRMC and S-MCANS on the ML-100K dataset.

7.1 Further discussion and related work

Bandit problems where multiple actions are selected have also been considered in the past. Kale et al. [2010] consider a setup where on choosing multiple arms the reward obtained is the sum of the rewards of the chosen arms, and the reward of each chosen arm is revealed to the algorithm. Both these works focus on obtaining guarantees on the

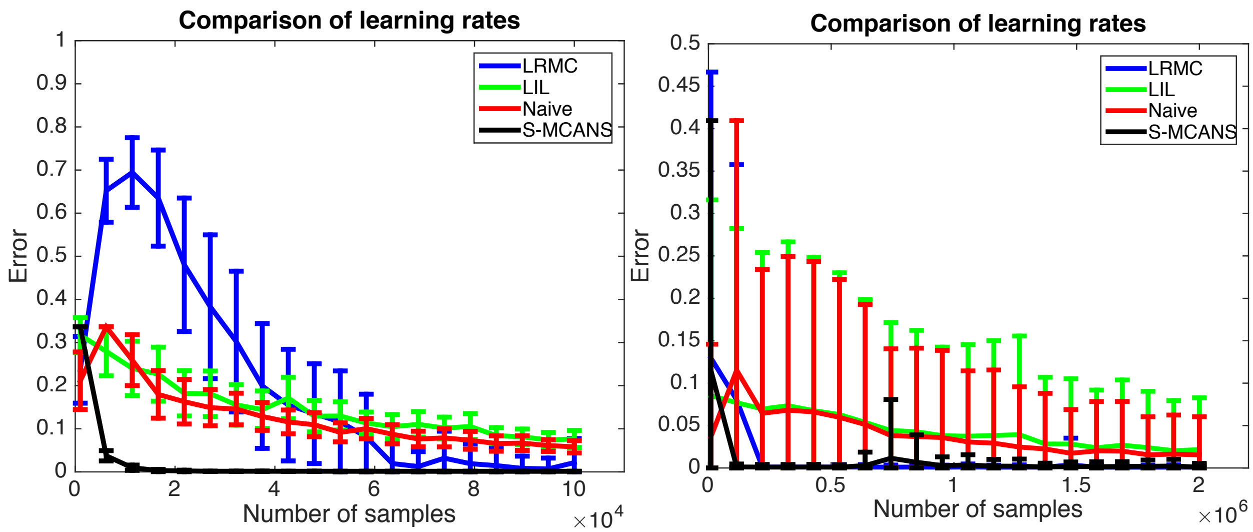

(a) Synthetic; $(80, 2)$

(b) Synthetic; $(400, 2)$

Figure 1: Error of various algorithms with increasing budget. Numbers in the brackets represent values for (K, r). The error is defined as $\boldsymbol{L}_{\hat{i},\hat{j}} - \boldsymbol{L}_{i_\star,j_\star}$ where $(\hat{i}, \hat{j})$ is a pair of optimal choices as estimated by each algorithm.

cumulative regret compared to the best set of arms in hindsight. Radlinski et al. [2008] consider a problem, in the context of information retrieval, where multiple bandit arms are chosen and the reward obtained is the maximum of the rewards corresponding to the chosen arms. Apart from this reward information the algorithm also gets a feedback that tells which one of the chosen arms has the highest reward. Similar models have also been studied in Streeter and Golovin [2009] and Yue and Guestrin [2011]. A major difference between the above mentioned works and our work is the feedback and reward model and the fact that we are not interested in regret guarantees but rather in finding a good pair of arms as quickly as possible. Furthermore our linear-algebraic approach to the problem is very different from previous approaches which were either based on multiplicative weights [Kale et al., 2010] or online greedy submodular maximization [Streeter and Golovin, 2009, Yue and Guestrin, 2011, Radlinski et al., 2008]. Simchowitz et al. [2016] also consider similar subset selection problems and provide algorithms to identify the top set of arms. In the Web search literature click models have been proposed to model user behaviour [Guo et al., 2009, Craswell et al., 2008] and a bandit analysis of such models have also been proposed [Kveton et al., 2015]. However, these models assume that all the users come from a single population and tend to use richer information in their formulations (for example information about which exact link was clicked). Finally we would like to mention that our model shown in Figure 5.1 of the main paper on the surface bears resemblance to dueling bandit problems [Yue et al., 2012]. However, in duleing bandits two arms are compared which is not the case in the bandit problem that we study.Interactive collaborative filtering (CF) and bandit approaches to such problems have also been investigated [Kawale et al., 2015]. Though, the end goal in CF is different from our goal in this paper.

References

N. Craswell, O. Zoeter, M. Taylor, and B. Ramsey. An experimental comparison of click position-bias models. In *WSDM*. ACM, 2008.

F. Guo, C. Liu, A. Kannan, T. Minka, M. Taylor, Y.-M. Wang, and C. Faloutsos. Click chain model in web search. In *WWW*. ACM, 2009.

S. Kale, L. Reyzin, and R. E. Schapire. Non-stochastic bandit slate problems. In *NIPS*, 2010.

J. Kawale, H. H. Bui, B. Kveton, L. Tran-Thanh, and S. Chawla. Efficient thompson sampling for online matrix-factorization recommendation. In *NIPS*, 2015.

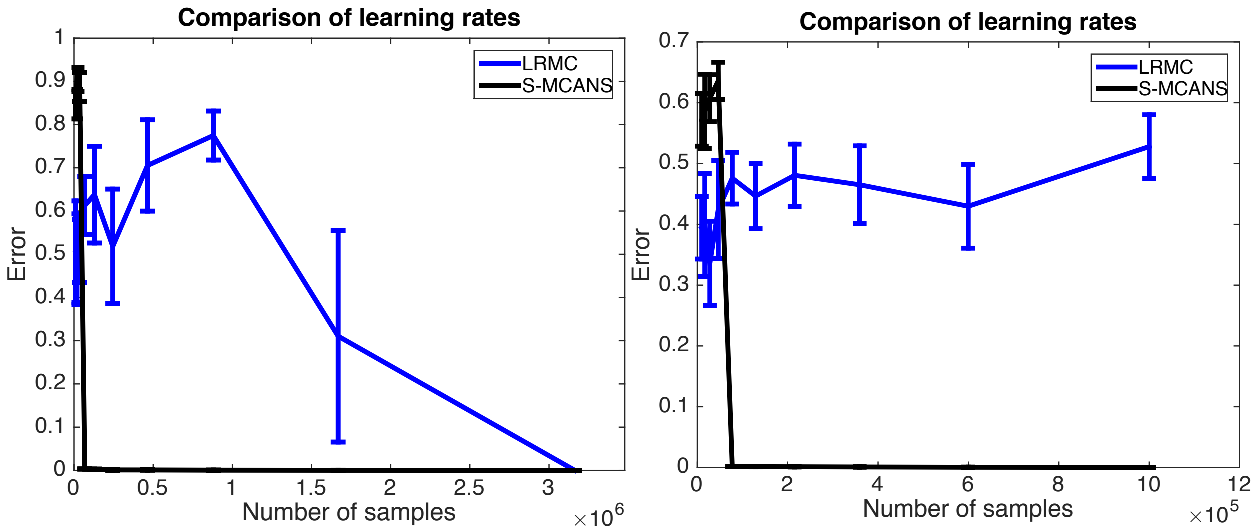

(a) ML-100K; $(200, 2)$ (b) ML-100K; $(200, 4)$

Figure 2: Error of LRMC and S-MCANS algorithms with increasing budget. Numbers in the brackets represent values for (K, r). The error is defined as $\boldsymbol{L}_{\hat{i},\hat{j}} - \boldsymbol{L}_{i_\star,j_\star}$ where $(\hat{i},\hat{j})$ is a pair of optimal choices as estimated by each algorithm.. This is for the ML-100K dataset

B. Kveton, C. Szepesvari, Z. Wen, and A. Ashkan. Cascading bandits: Learning to rank in the cascade model. In *ICML*, pages 767–776, 2015.

F. Radlinski, R. Kleinberg, and T. Joachims. Learning diverse rankings with multi-armed bandits. In *ICML*. ACM, 2008.

M. Simchowitz, K. Jamieson, and B. Recht. Best-of-k bandits. In *COLT*, 2016.

M. Streeter and D. Golovin. An online algorithm for maximizing submodular functions. In *NIPS*, 2009.

J. A. Tropp. An introduction to matrix concentration inequalities. *arXiv preprint arXiv:1501.01571*, 2015.

Y. Yue and C. Guestrin. Linear submodular bandits and their application to diversified retrieval. In *NIPS*, pages 2483–2491, 2011.

Y. Yue, J. Broder, R. Kleinberg, and T. Joachims. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.