Recommender Systems and Statistical Guarantee for Collaborative Filtering

Ni Zhan (nzhan)

Yilin Yang (yiliny2) **Zicheng Cai (zichengc)**

1 Introduction

Recommender systems are an area of interest for music, shopping, and other content. With popularity of e-commerce, digital music platforms, etc., personalized recommender systems have high commercial value. An effective recommender system should account for users' previous preferences and recommend items which the user is likely to enjoy or purchase. Some challenges are the large database of available items, high dimensionality and heterogeneity of data representing the items (songs, movies, etc.). Collaborative filtering is widely used in recommender systems, and uses feedback from many users on the items. Some collaborative filtering methods disregard the meta-data, while recent context-aware methods utilize meta-data, which are expected to boost performance. As example, meta-data of users could be their social network or demographics data, and meta-data of movies could be genre, year, director, awards, etc. We study theoretical guarantees of collaborative filtering with noisy observations and conduct an empirical study comparing it with related methods.

In collaborative filtering, the feedback matrix $\mathbf{Y} \in \mathbb{R}^{m \times n}$ with m users and n items contains feedback from each user for the items. The feedback could be explicit such as ratings, or implicit such as number of listens, number of clicks, views, etc. Entries of $\mathbf{Y}$ are partially observed, and the goal is to infer the missing entries. To pose this as a meaningful problem, we assume the variables lie in a lower-dimensional subspace, i.e. $\mathbf{Y}$ is low rank and can be written as $\mathbf{Y} = \mathbf{W}\mathbf{H}^T$ where $\mathbf{W} \in \mathbb{R}^{m \times k}$ is a low-rank representation of users, and $\mathbf{H} \in \mathbb{R}^{n \times k}$ represents items [2]. This can be solved with the regularized optimization

$$\min_{\mathbf{Z}} \|P_\Omega(\mathbf{Z} - \mathbf{Y})\|_F^2 + \lambda \text{rank}(\mathbf{Z})$$

where $P_\Omega(\cdot)$ is the projection operator that retains only observed entries [3]. Rank minimization is NP-hard, and we solve the convex relaxation substituting rank($\mathbf{Z}$) with $\|\mathbf{Z}\|_*$ the nuclear norm

$$\min_{\mathbf{Z}} \|P_\Omega(\mathbf{Z} - \mathbf{Y})\|_F^2 + \lambda \|\mathbf{Z}\|_*. \tag{1}$$

We will study problem (1) and its extensions as shown in [1, 4]. One natural motivation is including meta-data to the collaborative filtering problem to utilize more information and boost performance. Rao et al. included meta-data through graph structures where similar users or items are connected by a graph edge [6]. They assume graph (V^w, E^w) adjacency matrix encodes relationships between m rows of $\mathbf{W}$ and graph (V^h, E^h) for n rows of $\mathbf{H}$ [6]. To encourage $\mathbf{W}$ to follow its graph structure, the term tr($\mathbf{W}^T \mathbf{Lap}(E^w)\mathbf{W}$) is added to the minimization problem, where $\mathbf{Lap}(\cdot)$ is the graph Laplacian [7], and similarly for $\mathbf{H}$. We show the statistical guarantee for the resulting problem (2) is similar to the one for problem (1) by casting (2) as a weighted nuclear norm regularization.

$$\min_{\mathbf{W},\mathbf{H}} \|P_\Omega(\mathbf{W}\mathbf{H}^T - \mathbf{Y})\|_F^2 + \lambda_w \|\mathbf{W}\|_F^2 + \lambda_h \|\mathbf{H}\|_F^2 + \lambda_L \text{tr}(\mathbf{W}^T \mathbf{Lap}(E^w)\mathbf{W}) + \lambda_L \text{tr}(\mathbf{H}^T \mathbf{Lap}(E^h)\mathbf{H}). \tag{2}$$

Preprint. Under review.

2 Notation and Assumptions

2.1 Matrix Incoherence Assumptions

In general, it is not possible to solve problem (1) for any matrix. As an example, for a rank-2 matrix with all zeros except the top-left 2×2 corner, we would need to see basically all entries before recovering the matrix. Candès and Recht summarized this idea, that the left and right singular vectors of the matrix need to be uncorrelated with the standard basis, and introduced the property of coherence [1].

Definition 2.1 (Coherence)**.** Let U be a subspace of $\mathbb{R}^n$ of dimension k, P_U the orthogonal projection onto U, and $\mathbf{e}_i$ the standard basis. The coherence of U is defined as

$$\mu(U) := \frac{n}{k} \max_{1 \le i \le n} \|P_U \mathbf{e}_i\|^2. \tag{3}$$

The largest value of $\mu(U)$ is $\frac{n}{k}$ and achieved for a subspace that contains a standard basis element. We will see that the matrices which can be recovered in problem (1) have column and row spaces of low coherence.

For the first theoretical result and non-noisy observation model, we introduce some assumptions. Suppose $\mathbf{Y}$ has SVD $\mathbf{Y} = \sum_{j=1}^{k} \sigma_j \mathbf{u}_j \mathbf{v}_j^*$ with column and row spaces U and V, respectively.

Assumption 1. *The coherences are upper bounded.* $\max(\mu(U), \mu(V)) \le \mu_0$ *for some positive* μ_0

Assumption 2. *The* $m \times n$ *matrix* $\sum_{j=1}^{k} \mathbf{u}_j \mathbf{v}_j^*$ *has a maximum entry, bounded by* $\mu_1 \sqrt{r/(mn)}$ *in absolute value for some positive* μ_1*.*

2.2 Spikiness and Rank Measures

The assumptions 1 and 2 are somewhat restrictive, and Negahban and Wainwright developed alternate conditions to establish error bounds for problem (1) with noisy observations [4]. Here we introduce the noisy observation model that they used. Suppose there exists $Z^* \in \mathbb{R}^{m \times n}$ with rank k. N entries from Z^* are sampled uniformly and observed, which are $Y = P_\Omega(Z^*)$. We assume the observed samples have added noise according to

$$Y_{a(i),b(i)} = Z^*_{a(i),b(i)} + \frac{\sigma}{\sqrt{mn}} \eta_i \qquad \eta \sim \mathcal{N}(0, 1). \tag{4}$$

For convenience we define $X_i := \sqrt{mn} \epsilon_i \mathbf{e}_{a(i)} \mathbf{e}_{b(i)}^T$ where ϵ_i is a random sign $\{-1, +1\}$, $\mathscr{X}(Z) \in \mathbb{R}^N$ as the vector such that $\mathscr{X}(Z)_i := \langle X_i, Z \rangle$ for $i = 1, ..., N$, and $y \in R^N$ is the vector of corresponding observations. Now $y_i = \langle X_i, Z^* \rangle + \sigma\eta$ is statistically equivalent to model (4). For any matrix $Z \in R^{m \times n}$, $\mathbb{E}[\langle X_i, Z \rangle^2] = \sum_{j=1}^{m} \sum_{k=1}^{n} R_j Z_{jk}^2 C_k = \|\sqrt{R} Z \sqrt{C}\|_F^2$. We define $\|\sqrt{R} Z \sqrt{C}\|_F^2$ as the weighted Frobenius norm $\| \cdot \|_{\omega(F)}$ in terms of weights R and C.

To establish error bounds on the problem (1), we need additional conditions on Z. We introduce measurements of the matrix:

$$\alpha(Z) := \sqrt{mn} \frac{\|\sqrt{R} Z \sqrt{C}\|_\infty}{\|\sqrt{R} Z \sqrt{C}\|_F} := \sqrt{mn} \frac{\|Z\|_{\omega(\infty)}}{\|Z\|_{\omega(F)}},$$

$$\beta(Z) := \frac{\|\sqrt{R} Z \sqrt{C}\|_*}{\|\sqrt{R} Z \sqrt{C}\|_F} := \frac{\|Z\|_{\omega(*)}}{\|Z\|_{\omega(F)}}.$$

Intuitively, $\alpha(Z)$ is related to the "spikiness" of Z. For highly spiky matrices, such as the maximally spiky matrix that has only one nonzero entry, we cannot establish error bounds. In addition, $1 \le \alpha(Z) \le \sqrt{mn}$. $\beta(Z)$ is related to the actual rank of Z.

2.3 Extension to Graph-Based Regularization

We cast the regularizer of (2) as a generalized version of the weighted nuclear norm. The weights will be related to the graph Laplacians $\mathbf{Lap}(E^w), \mathbf{Lap}(E^h)$. We define $L_w := \lambda_L \mathbf{Lap}(E^w) + \lambda_w I$

and $L_h := \lambda_L \mathbf{Lap}(E^h) + \lambda_h I$. Now let $L_w = U_w S_w U_w^T$ and $L_h = U_h S_h U_h^T$ be eigendecompositions for L_w and L_h. The weighted norm will be related to $A = U_w S_w^{-1/2}$ and $B = U_h S_h^{-1/2}$. We define the graph-weighted norm $\|Z\|_{\mathscr{A}(*)} := \|A^{-1} Z B^{-T}\|_*$. The corresponding graph-based measurements are

$$\alpha_g(Z) := \sqrt{mn} \frac{\|A^{-1} Z B^{-T}\|_\infty}{\|A^{-1} Z B^{-T}\|_F}, \qquad \beta_g(Z) := \frac{\|A^{-1} Z B^{-T}\|_*}{\|A^{-1} Z B^{-T}\|_F}.$$

3 Key Results

The first result uses the matrix incoherence properties and noiseless observations.

Theorem 1. *We assume $\mathbf{Y} \in \mathbb{R}^{m \times n}$ has rank k and obeys Assumptions 1 and 2, and let $d = \max(m, n)$. We observe N entries sampled uniformly at random. For*

$$N \geq C \max(\mu_1^2, \mu_0^{1/2} \mu_1, \mu_0 d^{1/4}) dk(\beta \log d)$$

where $\beta > 2$, C is positive constant, the solution to problem (1) is unique and equal to $\mathbf{Y}$ with high probability.

Theorem 1 shows that the low-rank matrix can be recovered with sufficient observed samples. In low rank, low coherence cases, the number of samples is much smaller than the total entries of the matrix. For noisy samples and an alternative set of matrix assumptions using α and β, Theorem 2 states the main theoretical result.

Theorem 2. *We assume Z^* is at most rank k, $\|Z^*\|_F \leq 1$, and $\alpha(Z^*) \leq \alpha^*$. For the convex problem*

$$\hat{Z} = \arg\min_{Z \in \mathcal{C}} \frac{1}{2N} \|y - \mathscr{X}(Z)\|_2^2 + \lambda \|Z\|_{\omega(*)} \text{ where } \mathcal{C} := \left\{ Z : \|Z\|_{\omega(\infty)} \leq \frac{\alpha^*}{\sqrt{mn}} \right\}$$

with $\lambda = C_1 \sigma \sqrt{\frac{(m+n)\log(m+n)}{N}}$, with high probability,

$$\|\hat{Z} - Z^*\|_{\omega(F)}^2 \leq C \alpha^{*2} \max\{1, \sigma^2\} \frac{k(m+n)\log(m+n)}{N}$$

where C, C_1 are positive constants.

The set $\mathcal{C}$ restricts the matrices under consideration in terms of spikiness and rank measures. The restricted set ensures that the loss is strictly convex, and Theorem 2 shows we are able to recover the original matrix with high probability. Theorem 3 is the corresponding theorem for problem (2). We will show that (2) may be cast as generalized weighted nuclear norm regularization. Theorem 3 follows the same noisy observation model as Eq. 4.

Theorem 3. *We assume Z^* is low rank k, $\|Z^*\|_F = 1$, and the minimum singular values of A, B (as defined in Section 2.3) are 1. For the convex problem*

$$\hat{Z} = \arg\min_{Z \in \mathcal{C}} \frac{1}{N} \|y - \mathscr{X}(Z)\|_2^2 + \lambda \|Z\|_{\mathscr{A}(*)} \text{ where } \mathcal{C} := \left\{ Z : \alpha_g(Z)\beta_g(Z) \leq c_0 \sqrt{\frac{N}{\log(m+n)}} \right\}$$

with $\lambda \geq C_1 \sigma \sqrt{\frac{(m+n)\log(m+n)}{N}}$, with high probability,

$$\|\hat{Z} - Z^*\|_F^2 \leq C \alpha_g^{*2} \max\{1, \sigma^2\} \frac{k(m+n)\log(m+n)}{N}$$

where C, C_1, c_0 are positive constants, and $\alpha_g^* := \alpha_g(Z^*)$.

Rao et al. propse a conjugate gradient algorithm to solve the subproblems for W, H of (2). Authors of [1] and [4] used semi-definite programming (SDP) solvers for their problems.

4 Proof Sketches

4.1 Proof Sketch of Theorem 2

The proof requires restricted strong convexity (RSC) of the loss term $\mathcal{L} = \frac{1}{2N}\|y - \mathscr{X}(Z)\|_2^2$, and RSC is defined in [5]. When RSC holds, $\frac{\|\mathscr{X}(Z)\|_2}{\sqrt{N}} \geq c\|Z\|_{\omega(F)}$ for a positive constant c. This indicates that the Hessian $\nabla^2\mathcal{L}$ is strictly positive definite for the restricted set of Z. We state the following, which guarantees RSC in our case. For the constraint set

$$\mathscr{C} := \left\{ \Delta \in \mathbb{R}^{m\times n}, \Delta \neq 0 \quad | \quad \alpha(\Delta)\beta(\Delta) \leq c_0 \sqrt{\frac{N}{(m+n)\log(m+n)}} \right\}, \tag{5}$$

$$\frac{1}{\sqrt{N}}\|\mathscr{X}(\Delta)\|_2 \geq \frac{1}{8}\|\Delta\|_{\omega(F)} \left\{ 1 - c_1 \frac{\alpha(\Delta)}{\sqrt{N}} \right\} \qquad \forall \Delta \in \mathscr{C} \tag{6}$$

and c_0, c_1 are positive constants.

For convenience, define $\Gamma := \sqrt{R}Z\sqrt{C}$, and $\mathscr{X}'(\Gamma) := \mathscr{X}(Z)$. Note that $\mathscr{X}'(\Gamma)_i = \langle \tilde{X}_i, \Gamma\rangle$ where $\tilde{X}_i = R^{-1/2}X_iC^{-1/2}$. We have

$$\hat{\Gamma} = \arg \min_{\|\Gamma\|_\infty \leq \frac{\alpha^*}{\sqrt{mn}}} \frac{1}{2N}\|y - \mathscr{X}'(\Gamma)\|_2^2 + \lambda\|\Gamma\|_* \quad .$$

Our goal is to upper bound $\|\hat{\Gamma} - \Gamma^*\|_F$. We define $\hat{\Delta} := \hat{\Gamma} - \Gamma^*$, and split $\hat{\Delta}$ into cases.

4.1.1 Case 1

The first case is $\hat{\Delta} \notin \mathscr{C}$. By definition (5),

$$\begin{aligned}
\|\hat{\Delta}\|_F^2 &\leq c_0\sqrt{mn}\|\hat{\Delta}\|_\infty \|\hat{\Delta}\|_* \sqrt{\frac{(m+n)\log(m+n)}{N}} \\
&\leq 2c_0\alpha^*\|\hat{\Delta}\|_* \sqrt{\frac{(m+n)\log(m+n)}{N}}
\end{aligned}$$

by $\|\hat{\Delta}\|_\infty \leq \|\Gamma^*\|_\infty + \|\hat{\Gamma}\|_\infty \leq \frac{2\alpha^*}{\sqrt{mn}}$. To bound the nuclear norm, we decompose $\hat{\Delta} = \hat{\Delta}' + \hat{\Delta}''$, where $\hat{\Delta}'$ has rank at most $2k$, and the nuclear norm of $\hat{\Delta}''$ is upper bounded by nuclear norm of $\hat{\Delta}'$ plus sum of additional singular values. Combining with the triangle inequality, we obtain

$$\|\hat{\Delta}\|_F^2 \leq c\alpha^* \sqrt{\frac{(m+n)\log(m+n)}{N}} \{8\sqrt{k}\|\hat{\Delta}\|_F + 4\sum_{j=k+1}^{m} \sigma_j(\Gamma^*)\} \tag{7}$$

4.1.2 Case 2

Consider $\hat{\Delta} \in \mathscr{C}$, and $c_1\frac{\alpha(\hat{\Delta})}{\sqrt{N}} > 1/2$. Substituting for $\alpha(\hat{\Delta})$, we have

$$\|\hat{\Delta}\|_F < \frac{c\sqrt{mn}\|\hat{\Delta}\|_\infty}{\sqrt{N}} \leq 2c\frac{\alpha^*}{\sqrt{N}} \tag{8}$$

where we again used $\|\hat{\Delta}\|_\infty \leq \frac{2\alpha^*}{\sqrt{mn}}$, for the second inequality.

On the other hand, consider $\hat{\Delta} \in \mathscr{C}$, and $c_1\frac{\alpha(\hat{\Delta})}{\sqrt{N}} \leq 1/2$. From condition (6), we have

$$\frac{\|\mathscr{X}'(\hat{\Delta})\|_2}{\sqrt{N}} \geq \frac{1}{16}\|\hat{\Delta}\|_F. \tag{9}$$

Because $\hat{\Gamma}$ is optimal and Γ^* is feasible,

$$\frac{1}{2N}\|y - \mathscr{X}'(\hat{\Gamma})\|_2^2 + \lambda\|\hat{\Gamma}\|_* \leq \mathscr{X}'(\Gamma^*)\|_2^2 + \lambda\|\Gamma^*\|_*$$

After some algebra yields

$$\frac{1}{2N}\|\mathscr{X}'(\hat{\Delta})\|_2^2 \leq \sigma\langle\hat{\Delta}, \frac{1}{N}\sum_{i=1}^{N}\epsilon_i\eta_i\tilde{X}_i\rangle + \lambda\|\Gamma^*\|_* - \lambda\|\Gamma^* + \hat{\Delta}\|.$$

By Holder's inequality, the triangle inequality, and substituting (9), we have

$$\|\hat{\Delta}\|_F^2 \leq c\sigma\|\frac{1}{N}\sum_{i=1}^{N}\epsilon_i\eta_i\tilde{X}_i\|_{op}\|\hat{\Delta}\|_* + \lambda\|\hat{\Delta}\|_* \tag{10}$$

where $\|\cdot\|_{op}$ indicates the spectral norm. We use the same upper bound for $\|\hat{\Delta}\|_*$ as in Section 4.1.1, and this yields

$$\|\hat{\Delta}\|_F^2 \leq c(\sqrt{k}\|\hat{\Delta}\|_F + \sum_{j=k+1}^{m}\sigma_j(\Gamma^*))(\sigma\|\frac{1}{N}\sum_{i=1}^{N}\epsilon_i\eta_i\tilde{X}_i\|_{op} + \lambda). \tag{11}$$

Because Z^* is at most rank k by assumption, the terms $\sum_{j=k+1}^{m}\sigma_j(\Gamma^*) = 0$ and can be removed from conditions (8) and (11). By the Ahlswede-Winter matrix bound, it can be shown that $\|\frac{1}{N}\sum_{i=1}^{N}\epsilon_i\eta_i\tilde{X}_i\|_{op} \geq \max\{c, \sqrt{\frac{(m+n)\log(m+n)}{N}}\}$. Combining the pieces, we need bound (7), (8), or (11) to hold. We substitute λ as specified in Theorem 2 into bound (11). Since $\alpha \geq 1$ always, we can summarize the conditions together as

$$\|\hat{\Delta}\|_F^2 = \|\hat{Z} - Z^*\|_{\omega(F)}^2 \leq C\alpha^{*2}\max\{1, \sigma^2\}\frac{k(m+n)\log(m+n)}{N}.$$

4.2 Proof Sketch of Theorem 3

We first need to establish equivalence between the weighted nuclear norm and graph regularization. We want to show $\|Z\|_{\mathscr{A}(*)} := \|A^{-1}ZB^{-T}\|_* = \inf_{W,H}\frac{1}{2}(\|A^{-1}W\|_F^2 + \|B^{-1}H\|_F^2)$. If this is true, then $\mathrm{tr}(W^TL_wW) = \|A^{-1}W\|_F^2 = \|S_w^{\frac{1}{2}}U_w^{-1}W\|_F^2$ and similarly, $\mathrm{tr}(H^TL_hH) = \|B^{-1}H\|_F^2 = \|S_h^{\frac{1}{2}}U_h^{-1}H\|_F^2$, by our definitions in Section 2.3. It would then be clear that problem (2) is a case of the optimization problem of Theorem 3.

Z has a SVD and can be written as $Z = \sum_i c_iA\mathbf{u}_i\mathbf{v}_i^TB^T$ with $\|\mathbf{u}_i\| = \|\mathbf{v}_i\| = 1$. It follows that $\|A^{-1}ZB^{-T}\|_* = \sum_i|c_i|$. We also can write $Z = WH^T$. In this case, the i-th column of W, H are $\mathbf{w}_i = \sqrt{|c_i|}A\mathbf{u}_i$ and $\mathbf{h}_i = \sqrt{|c_i|}A\mathbf{v}_i$. Now, $\|A^{-1}W\|_F^2 = \|B^{-1}H\|_F^2 = \sum_i|c_i|$, which shows $\|A^{-1}ZB^{-T}\|_* \geq \inf_{W,H}\frac{1}{2}(\|A^{-1}W\|_F^2 + \|B^{-1}H\|_F^2)$. For the other side, construct $\mathbf{u}_i = \frac{A^{-1}\mathbf{w}_i}{\|A^{-1}\mathbf{w}_i\|}$, $\mathbf{v}_i = \frac{B^{-1}\mathbf{h}_i}{\|B^{-1}\mathbf{h}_i\|}$, and $c_i = \|A^{-1}\mathbf{w}_i\|\|B^{-1}\mathbf{h}_i\|$. We have $Z = \sum_i c_iA\mathbf{u}_i\mathbf{v}_i^TB^T = \sum_i\mathbf{w}_i\mathbf{h}_i^T$. Now, $|c_i| = \|A^{-1}\mathbf{w}_i\|\|B^{-1}\mathbf{h}_i\| \leq \frac{1}{2}(\|A^{-1}\mathbf{w}_i\|^2 + \|B^{-1}\mathbf{h}_i\|^2)$ by the inequality of arithmetic and geometric means. This shows $\|A^{-1}ZB^{-T}\|_* \leq \inf_{W,H}\frac{1}{2}(\|A^{-1}W\|_F^2 + \|B^{-1}H\|_F^2)$.

The remainder of the proof could be similar to Section 4.1, however Rao et al. use a slightly different approach. Specifically they use Theorem 1 of [5], which is stated here

Theorem 4. *For the convex optimization problem*

$$\hat{Z} = \arg\min_{Z\in\mathbb{R}^{m\times n}} = (Z; y) + \lambda\mathcal{R}(Z)$$

where the regularizer $\mathcal{R}$ is a norm and decomposable w.r.t. the subspace pair $(\mathcal{M}, \mathcal{M}^\perp)$ and $\mathcal{M} \subseteq \mathbb{R}^{m\times n}$; the loss function $\mathcal{L}$ is convex and differentiable, and satisfies restricted strong convexity with curvature κ; and λ is a strictly positive constant $\lambda \geq 2\mathcal{R}^*(\nabla\mathcal{L}(Z^*))$, for $Z^* \in \mathcal{M}$, we have

$$\|\hat{Z} - Z^*\|^2 \leq 9\frac{\lambda^2}{\kappa^2}\Psi(\mathcal{M})^2$$

where $\Psi(\mathcal{M}) := \sup_{Z\in\mathcal{M}\setminus\{0\}}\frac{\mathcal{R}(Z)}{\|Z\|_F}$.

In the steps of the proof, they show that $\|Z\|_{\mathscr{A}(*)} \leq \sqrt{k}$, and with the assumption $\|Z\|_F = 1$, it follows that $\Psi(\mathcal{M}) \leq \sqrt{k}$. They also show $\mathcal{R}^*(Z) = \|Z\|^*_{\mathscr{A}(*)} = \|A^T Z B\|$, and find an upper bound for $R^*(\nabla \mathcal{L}(Z^*)) \leq C\sigma\sqrt{\frac{(m+n)\log(m+n)}{N}}$. This upper bound is set equal to λ as required by Theorem 4. Finally, they show the restricted strong convexity holds, and it is similar to Eq. 6 for the specified set $\mathcal{C}$ of Theorem 3. From [5], κ is related to the lower bound of the eigenvalues of the Hessian $\nabla^2 \mathcal{L}(Z^*)$, and the RSC condition indicates that $\kappa \propto 1/\alpha_g^*$.

5 Experimental Results

To corroborate Theorem 2, we performed a experiment to verify the upper bound provided in Theorem 2 and investigate the effect of the number of obervations on the total loss (the square of the Frobenius norm of the true matrix and the predicted matrix). According to the assumption of Theorem 2, we used a synthetic random matrix with 20 rows and 20 columns. We contrained the rank of the matrix to be less than or equal to 5 while the Frobenius norm of the matrix is less than or equal to 1. These are the basic two requirements for the matrix. Specifically, to generate a matrix M of m rows and n columns, we first generate two matrices U and V with the shape of $m \times k$ and $n \times k$ respectively. Then we randomly masked out the entries of these two matrices with a certain probability. These two steps help to get a matrix M with low rank and small norm from UV^T.

Before applying the matrix in the experiment, we double checked its rank and norm and ensure it satisfies the requirements. After that, we randomly draw N samples from the generated matrix as the training set. For simplicity, we assume the samples are distributed uniformly across the rows and the columns. Thus, the matrices A and B when calculating the α and β in 2 are both identity matrix. For the training part, we used the constrained optimization module of scipy to obtain the solution that minimizez the MSE loss and the nuclear norm loss while keeping small infinity norm as shown in Theorem 2. Since there is some randomness during the experiments (e.g., matrix generation, observation selection), we repeated the experiment with the same parameters (e.g., size of the matrix, number of the observations) for 10 times.

The average and the standard deviation of the results are reported in Figure 1 and Figure 2. Figure 1 shows the provided upper bound of the SSE from Theorem 2. Clearly, the upper bound decreases with more observation provided to the training set. Figure 2 illustrates the trend of the SSE with the number of observations. Similar to the upper bound, it also decreases along with the number of observations. We note that the upper bound has much greater magnitude compared to the SSE. This might be attributed to the choice of the positive constants C and C_1 in Theorem 2. We did not find the ways to estimate these two constants in the paper, thus we arbitrarily set them as 1 in our calculation. Using smaller constants (e.g, set $C = 0.001$) will certainly lead to lower upper bound. Thus, how to estimate the value of these constants could be a future direaction to makes Theorem 2 more concrete.

In addtion, we investigated the effect of the rank of the target matrix on the theoretical upper bound and the experimental prediction error. The experimental parameters are the same as the above description, except in this experiment, we fixed the number of observations to be 50, and varied the rank of the target matrix from 1 to 5. Corresponding results are shown in Figure 3 and Figure 4. Apparently, both the upper bound and the prediction error increase with the rank of the target matrix. Intuitively, when we fix the size of the training set, the matrix with more information (higher rank) is more difficult to recover than the matrix with less information (lower rank).

6 Conclusion

In this project, we investigate the theoretical properties of the matrix completion provided in Theorem 1 to Theorem 3. These theorems construct the upper bound of the matrix recovery error under some assumptions. We discussed the theoretical proofs for these theorems in this project. In addition, we performed several experiments to verify the result of Theorem 2. In particular, we studied the effect of the obvervation size and the rank of the target matrix on the theoretical upper bound and the experimental matrix recovery error. We found that increasing the number of observations could result in lower error as expected and it is aligned with the theorems. Given a number of observations, matrices with higher ranks are harder to recovery in our experiment. This result is also reasonable

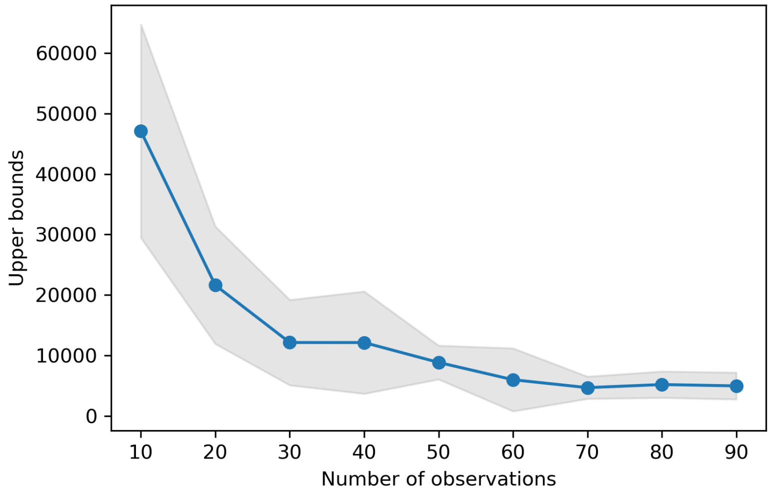

Figure 1: Upper bounds of the error provided by Theorem 2. The shaded region represents the standard deviation.

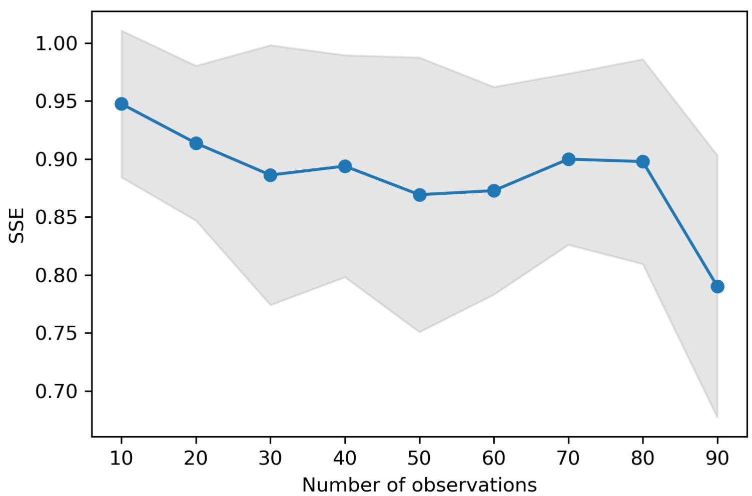

Figure 2: Sum of the square error with different observations. The shaded region represents the standard deviation.

according to Theorem 2. One of the issues that remain in our experiment is the estimation of two positive constants in Theorem 2. These two contants play an important role on determining the magnitude of the upper bound. Arbitrary setting may affect the usefulness of the upper bound. Thus, how to estimate these constants could be a future direction.

Division of Work

Ni Zhan: Theoretical proof sketch for the theorems. Zicheng Cai: Literature reviews. Yilin Yang: Experiment implementation.

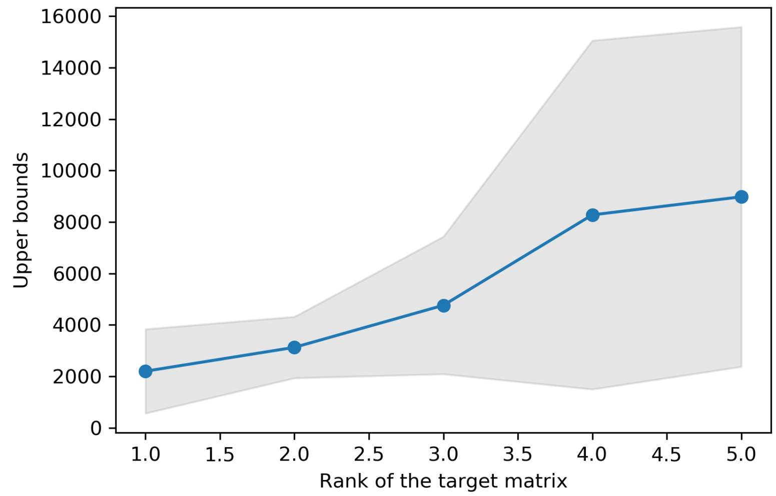

Figure 3: Upper bounds of the error provided by Theorem 2 with different ranks. The shaded region represents the standard deviation.

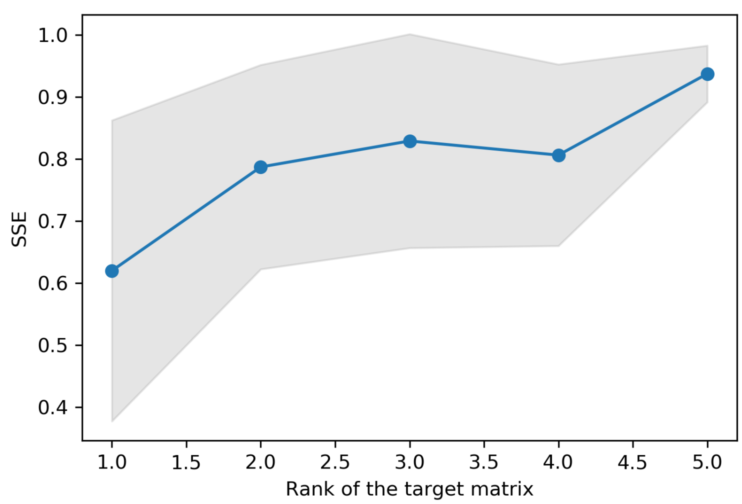

Figure 4: Sum of the square error with different ranks. The shaded region represents the standard deviation.

References

[1] Emmanuel J. Candès and Benjamin Recht. "Exact Matrix Completion Via Convex Optimization". In: *Foundations of Computational Mathematics* 9.6 (2009), pp. 717–772. DOI: 10.1007/s10208-009-9045-5. URL: https://doi.org/10.1007/s10208-009-9045-5.

[2] Brian McFee et al. "The million song dataset challenge". In: *Proceedings of the 21st International Conference on World Wide Web*. 2012, pp. 909–916.

[3] Federico Monti, Michael M. Bronstein, and Xavier Bresson. "Geometric Matrix Completion With Recurrent Multi-Graph Neural Networks". In: *CoRR* (2017). arXiv: 1704.06803 [cs.LG]. URL: http://arxiv.org/abs/1704.06803v1.

[4] Sahand Negahban and Martin J. Wainwright. "Restricted Strong Convexity and Weighted Matrix Completion: Optimal Bounds with Noise". In: *Journal of Machine Learning Research* 13.53 (2012), pp. 1665–1697. URL: http://jmlr.org/papers/v13/negahban12a.html.

[5] Sahand N. Negahban et al. "A Unified Framework for High-Dimensional Analysis of M-Estimators With Decomposable Regularizers". In: *Statistical Science* 27.4 (2012), nil. DOI: 10.1214/12-sts400. URL: https://doi.org/10.1214/12-sts400.

[6] Nikhil Rao et al. "Collaborative Filtering with Graph Information: Consistency and Scalable Methods". In: *Advances in Neural Information Processing Systems*. Ed. by C. Cortes et al. Vol. 28. Curran Associates, Inc., 2015. URL: https://proceedings.neurips.cc/paper/2015/file/f4573fc71c731d5c362f0d7860945b88-Paper.pdf.

[7] Eric W. Weisstein. *Laplacian Matrix. From MathWorld—A Wolfram Web Resource*. URL: https://mathworld.wolfram.com/LaplacianMatrix.html.