## Submitted to Operations Research manuscript (Please, provide the manuscript number!)

Authors are encouraged to submit new papers to INFORMS journals by means of a style file template, which includes the journal title. However, use of a template does not certify that the paper has been accepted for publication in the named journal. INFORMS journal templates are for the exclusive purpose of submitting to an INFORMS journal and should not be used to distribute the papers in print or online or to submit the papers to another publication.

# Interpretable Matrix Completion: A Discrete Optimization Approach 

Dimitris Bertsimas<br>Sloan School of Management and Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, dbertsim@mit.edu<br>Michael Lingzhi Li<br>Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA 02139, mlli@mit.edu


#### Abstract

We consider the problem of matrix completion on an $n \times m$ matrix. We introduce the problem of Interpretable Matrix Completion that aims to provide meaningful insights for the low-rank matrix using side information. We show that the problem can be reformulated as a binary convex optimization problem. We design OptComplete, based on a novel concept of stochastic cutting planes to enable efficient scaling of the algorithm up to matrices of sizes $n=10^{6}$ and $m=10^{6}$. We report experiments on both synthetic and realworld datasets that show that OptComplete has favorable scaling behavior and accuracy when compared with state-of-the-art methods for other types of matrix completion, while providing insight on the factors that affect the matrix.


Key words: Matrix Completion, Mixed-Integer Optimization, Stochastic Approximation

## 1. Introduction

Low-rank matrix completion has attracted much attention after the successful application in the Netflix Competition. It is now widely utilized in far-reaching areas such as computer vision (Candes and Plan (2010)), signal processing (Ji et al. (2010)), and control theory (Boyd et al. (1994)) to generate a completed matrix from partially observed entries.

The classical low-rank matrix completion problem considers the following problem: Given a matrix $\boldsymbol{A} \in \mathbb{R}^{n \times m}$ with entries only partially known (denote $\Omega \subset\{1, \ldots, n\} \times\{1, \ldots, m\}$ as the set of known entries), we aim to recover a matrix $\boldsymbol{X} \in \mathbb{R}^{n \times m}$ of rank $k$ that minimizes a certain distance metric between $\boldsymbol{X}$ and $\boldsymbol{A}$ on the known entries of $\boldsymbol{A}$ :

$$
\min _{\boldsymbol{X}} \frac{1}{n m} \sum_{(i, j) \in \Omega}\left\|X_{i j}-A_{i j}\right\| \quad \text { subject to } \quad \operatorname{Rank}(\boldsymbol{X})=k,
$$

where we normalized the objective so that it is $O(1)$. The rank $k$ constraint on $\boldsymbol{X}$ can be equivalently formulated as the existence of two matrices $\boldsymbol{U} \in \mathbb{R}^{n \times k}, \boldsymbol{V} \in \mathbb{R}^{m \times k}$ such that $\boldsymbol{X}=\boldsymbol{U} \boldsymbol{V}^{T}$. Therefore, the problem can be restated as:

$$
\begin{equation*}
\min _{\boldsymbol{U}} \min _{\boldsymbol{V}} \frac{1}{n m} \sum_{(i, j) \in \Omega}\left\|X_{i j}-A_{i j}\right\| \quad \text { subject to } \quad \boldsymbol{X}=\boldsymbol{U} \boldsymbol{V}^{T} \tag{1}
\end{equation*}
$$

In many applications for matrix completion, it is customary for each row of the data to represent an individual and each column a product or an item of interest, and $A_{i j}$ being the response data of individual $i$ on item $j$. Therefore, the matrices $\boldsymbol{U}$ and $\boldsymbol{V}$ are commonly interpreted as the "user matrix" and "product matrix" respectively.

Let us denote each row of $\boldsymbol{U}$ as $\boldsymbol{u}_{i}$ and each column as $\boldsymbol{u}^{i}$ (similarly for $\boldsymbol{V}$ ). Then, $\boldsymbol{u}^{i}\left(\boldsymbol{v}^{i}\right)$ represents a "latent feature" for users (products), and in total there are $k$ latent features for users (products). The goal of matrix completion is thus to discover such latent features of the users and the products, so that the dot product of such features on user $i$ and product $j, X_{i j}=\boldsymbol{u}_{i} \cdot \boldsymbol{v}_{j}$, is the intended response of user $i$ on product $j$.

While this interpretation is intuitive, it does not offer insight on what the latent features of users and products mean. Inductive Matrix Completion, first considered in Dhillon et al. (2013), aims to rectify such problem by asserting that each of the $k$ latent features is a linear combination of $p>k$ known features. We focus on the one-sided information case, where only one of the user/product matrix is subject to such constraint. This case is more relevant as features related to users are often fragmented and increasingly constrained by data privacy regulations, while features about
products are easy to obtain. We would also have a short discussion later on why the two-sided information case is not interesting under the context of this paper.

Without loss of generality, we would assume the information is on the product matrix and denote the known feature matrix with $p$ features as $\boldsymbol{B} \in \mathbb{R}^{m \times p}$. As an example, if the items are movies, then $\boldsymbol{b}^{j}$ represents feature $j$ for a movie (actors, budget, running time, etc), and each product feature $\boldsymbol{v}^{j}$ needs to be linear combination of such features. Mathematically, this translates to the constraint:

$$
V=B S
$$

where $\boldsymbol{S} \in \mathbb{R}^{p \times k}$. Therefore, the inductive version of the problem in (1) can be written as:

$$
\begin{equation*}
\min _{U} \min _{S} \frac{1}{n m} \sum_{(i, j) \in \Omega}\left\|X_{i j}-A_{i j}\right\| \quad \text { subject to } \quad \boldsymbol{X}=\boldsymbol{U} \boldsymbol{S}^{T} \boldsymbol{B}^{T} . \tag{2}
\end{equation*}
$$

Although the inductive version of the problem adds more interpretability to the product latent features, they are still far from fully interpretable. For example, if we take the items to be movies, and features to be running time, budget, box office, and number of top 100 actors, then the generated features could look like:

$$
\begin{aligned}
& 5.6 \times \text { running time }-0.00067 \times \text { budget }+12 \times \# \text { of Top } 100 \text { actors, } \\
& 0.25 \times \text { box office }-5 \times \# \text { of Top } 100 \text { actors. }
\end{aligned}
$$

These features, although a linear combination of interpretable features, are not very interpretable itself due to the involvement of multiple factors with different units. Therefore, it cannot significantly help decision makers to understand "what is important" about the product. Furthermore, the appearance of the same factor in multiple features with different signs (as shown above) further complicates any attempt at understanding the result.

Therefore, we argue that instead of supposing the $k$ features are linear combinations of the $p$ known features, we should assume that the $k$ features are selected from the $p$ known features. This
formulation alleviates the two previous problems mentioned: it guarantees the latent features to be interpretable (as long as the original features are), and it prevents any duplicating features in the selected $k$ latent features. We denote this Interpretable Matrix Completion. We use the term interpretable, as opposed to inductive, to highlight that our approach, like sparse linear regression, gives actionable insights on what are the important features of matrix $\boldsymbol{A}$. We note that Interpretable Matrix Completion is considerably harder than inductive or the classical matrix completion problem as it is a discrete problem and selecting $k$ out of $p$ factors is exponential in complexity.

In this paper, we show that the Interpretable Matrix Completion problem can be written as a mixed integer convex optimization problem. Inspired by Bertsimas and van Parys (2020) for sparse linear regression, we reformulate the interpretable matrix completion problem as a binary convex optimization problem. Then we introduce a new algorithm OptComplete, based on stochastic cutting planes, to enable scalability for matrices of sizes on the order of $(n, m)=\left(10^{6}, 10^{6}\right)$. In addition, we provide empirical evidence on both synthetic and real-world data that OptComplete is able to match or exceed current state-of-the-art methods for inductive and general matrix completion on both speed and accuracy, despite OptComplete solving a more difficult problem.

Specifically, our contributions in this paper are as follows:

1. We introduce the interpretable matrix completion problem, and reformulate it as a binary convex optimization problem that can be solved using cutting planes methods.
2. We propose a new novel approach to cutting planes by introducing stochastic cutting planes. We prove that the new algorithm converges to an optimal solution of the interpretable matrix completion problem with exponentially vanishing failure probability.
3. We present computational results on both synthetic and real datasets that show that the algorithm matches or outperforms current state-of-the-art methods in terms of both scalability and accuracy.

The structure of the paper is as follows. In Section 2, we introduce the binary convex reformulation of the low-rank interpretable matrix completion problem, and how it can be solved through a cutting plane algorithm, which we denote CutPlanes. In Section 3, we introduce OptComplete, a stochastic cutting planes method designed to scale the CutPlanes algorithm in Section2, and show that it recovers the optimal solution of CutPlanes with exponentially vanishing failure probability. In Section 4, we report on computational experiments with synthetic data that compare OptComplete to Inductive Matrix Completion (IMC) introduced in Natarajan and Dhillon (2014) and SoftImpute-ALS (SIALS) by Hastie et al. (2015), two state-of-the-art matrix completion algorithms for inductive and general completion. We also compare OptComplete to CutPlanes to demonstrate the 20x to 60 x speedup of the stochastic algorithm. In Section , we report computational experiments on the Netflix Prize dataset. In Section 6 we provide our conclusions.

## Literature

Matrix completion has been applied successfully to many tasks, including recommender systems Koren et al. (2009), social network analysis Chiang et al. (2014) and clustering Chen et al. (2014b). After Candès and Tao (2010) proved a theoretical guarantee for the retrieval of the exact matrix under the nuclear norm convex relaxation, a lot of methods have focused on the nuclear norm problem (see Mazumder et al. (2010), Beck and Teboulle (2009), Jain et al. (2010), and Tanner and Wei (2013) for examples). Alternative methods include alternating projections by Recht and Ré (2013) and Grassmann manifold optimization by Keshavan et al. (2009). There has also been work where the uniform distributional assumptions required by the theoretical guarantees are violated, such as Negahban and Wainwright (2012) and Chen et al. (2014a).

Interest in inductive matrix completion intensified after Xu et al. (2013) showed that given predictive side information, one only needs $O(\log n)$ samples to retrieve the full matrix. Thus, most of this work (see Xu et al. (2013), Jain and Dhillon (2013), Farhat et al. (2013),

Natarajan and Dhillon (2014)) have focused on the case in which the side information is assumed to be perfectly predictive so that the theoretical bound of $O(\log n)$ sample complexity Xu et al. (2013) can be achieved. Chiang et al. (2015) explored the case in which the side information is corrupted with noise, while Shah et al. (2017) and Si et al. (2016) incorporated nonlinear combination of factors into the side information. Surprisingly, as pointed out by a recent article Nazarov et al. (2018), there is a considerable lack of effort to introduce sparsity/interpretability into inductive matrix completion, with Lu et al. (2016), Soni et al. (2016) and Nazarov et al. (2018) being among the only works that attempt to do so. Our work differs from the previous attempts in that previous attempts mainly focus on choosing latent features which are sparse linear combinations of the given features. In contrast interpretable matrix completion is aimed to select exactly $k$ features from the known features.

## 2. Interpretable Matrix Completion

In this section, we present the mathematical formulation of Interpretable Matrix Completion and how it can be reformulated as a binary convex problem that is based on Bertsimas and van Parys (2020). We show how this naturally leads to a cutting plane algorithm, and discuss its computational complexity. We also discuss the two-sided information case, and how that reduces to the sparse regression problem.

### 2.1. Binary Convex Reformulation of Interpretable Matrix Completion

The (one-sided) interpretable matrix completion problem can be written as a mixed binary optimization problem:

$$
\min _{\boldsymbol{U}} \min _{s \in S_{k}^{p}} \frac{1}{n m} \sum_{(i, j) \in \Omega}\left\|X_{i j}-A_{i j}\right\| \quad \text { subject to } \quad \boldsymbol{X}=\boldsymbol{U} \boldsymbol{S} \boldsymbol{B}^{T}
$$

where $\boldsymbol{S}=\operatorname{Diag}\left\{s_{1}, \ldots, s_{p}\right\} \in \mathbb{R}^{p \times p}$ and:

$$
S_{k}^{p}=\left\{s=\left(s_{1}, \ldots, s_{p}\right)^{T} \in\{0,1\}^{p}: \sum_{i=1}^{p} s_{i}=k\right\} .
$$

We note that given that $\sum_{i=1}^{p} s_{i}=k$, the rank of matrix $\boldsymbol{X}$ is indeed $k$. We further note that the coefficients of $\boldsymbol{S}$ can be taken to be binary without loss of generality, since if they are not and $\boldsymbol{S}=\operatorname{Diag}\left(1 / d_{1}, \ldots, 1 / d_{p}\right)$, then by applying the transformation:

$$
\begin{equation*}
\boldsymbol{U} \rightarrow \boldsymbol{U} \boldsymbol{D} \quad \boldsymbol{S} \rightarrow \boldsymbol{S} \boldsymbol{D}^{-1} \tag{3}
\end{equation*}
$$

for $\boldsymbol{D}=\operatorname{Diag}\left(d_{1}, \ldots, d_{p}\right)$, results in an equivalent problem with the coefficients of $\boldsymbol{S}$ being binary.

For this paper, we consider the squared norm, and for robustness purposes (see Bertsimas and van Parys (2020) and Bertsimas and Copenhaver (2018)), we add a Tikhonov regularization term to the original problem. Specifically, the (one-sided) interpretable matrix completion problem with regularization we address is

$$
\begin{equation*}
\min _{\boldsymbol{U}} \min _{s \in S_{k}^{p}} \frac{1}{n m}\left(\sum_{(i, j) \in \Omega}\left(X_{i j}-A_{i j}\right)^{2}+\frac{1}{\gamma}\|\boldsymbol{U}\|_{2}^{2}\right) \quad \text { subject to } \quad \boldsymbol{X}=\boldsymbol{U} \boldsymbol{S} \boldsymbol{B}^{T} . \tag{4}
\end{equation*}
$$

In this section, we show how that problem (4) can be reformulated as a binary convex optimization problem, and can be solved to optimality using a cutting plane algorithm. The main theorem and proof is presented below:

Theorem 1 Problem (4) can be reformulated as a binary convex optimization problem:

$$
\min _{\boldsymbol{s} \in S_{k}^{p}} c(\boldsymbol{s})=\frac{1}{n m} \sum_{i=1}^{n} \overline{\boldsymbol{a}}_{i}\left(\boldsymbol{I}_{m}+\gamma \boldsymbol{W}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{K}_{j}\right) \boldsymbol{W}_{i}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T},
$$

where $\boldsymbol{W}_{1}, \ldots, \boldsymbol{W}_{n} \in \mathbb{R}^{m \times m}$ are diagonal matrices:

$$
\left(\boldsymbol{W}_{i}\right)_{j j}= \begin{cases}1, & (i, j) \in \Omega \\ 0, & (i, j) \notin \Omega\end{cases}
$$

$\overline{\boldsymbol{a}}_{i}=\boldsymbol{W}_{i} \boldsymbol{a}_{i}, i=1 \ldots, n$, where $\boldsymbol{a}_{i} \in \mathbb{R}^{1 \times m}$ is the ith row of $\boldsymbol{A}$ with unknown entries taken to be 0 , and $\boldsymbol{K}_{j}=\boldsymbol{b}^{j}\left(\boldsymbol{b}^{j}\right)^{T} \in \mathbb{R}^{m \times m}, j=1, \ldots, p$ with $\boldsymbol{b}^{j} \in \mathbb{R}^{m \times 1}$ the $j$ th column of $\boldsymbol{B}$.

Proof: With the diagonal matrices $\boldsymbol{W}_{i}$ defined above, we can rewrite the sum in (4) over known entries of $\boldsymbol{A}, \sum_{(i, j) \in \Omega}\left(X_{i j}-A_{i j}\right)^{2}$, as a sum over the rows of $\boldsymbol{A}$ :

$$
\sum_{i=1}^{n}\left\|\left(\boldsymbol{x}_{i}-\boldsymbol{a}_{i}\right) \boldsymbol{W}_{i}\right\|_{2}^{2}
$$

where $\boldsymbol{x}_{i} \in \mathbb{R}^{1 \times m}$ is the $i$ th row of $\boldsymbol{X}$. Using $\boldsymbol{X}=\boldsymbol{U} \boldsymbol{S} \boldsymbol{B}^{T}$, then $\boldsymbol{x}_{i}=\boldsymbol{u}_{i} \boldsymbol{S} \boldsymbol{B}^{T}$ where $\boldsymbol{u}_{i} \in \mathbb{R}^{1 \times m}$ is the $i$ th row of $U$. Moreover,

$$
\|\boldsymbol{U}\|_{2}^{2}=\sum_{i=1}^{n}\left\|\boldsymbol{u}_{i}\right\|_{2}^{2} .
$$

Then, Problem (4) becomes:

$$
\min _{s \in S_{k}^{p}} \min _{U} \frac{1}{n m}\left(\sum_{i=1}^{n}\left(\left\|\left(\boldsymbol{u}_{i} \boldsymbol{S} \boldsymbol{B}^{T}-\boldsymbol{a}_{i}\right) \boldsymbol{W}_{i}\right\|_{2}^{2}+\frac{1}{\gamma}\left\|\boldsymbol{u}_{i}\right\|_{2}^{2}\right)\right) .
$$

We then notice that within the sum $\sum_{i=1}^{n}$ each row of $\boldsymbol{U}$ can be optimized separately, leading to:

$$
\begin{equation*}
\min _{\boldsymbol{s} \in S_{k}^{p}} \frac{1}{n m}\left(\sum_{i=1}^{n} \min _{\boldsymbol{u}_{i}}\left(\left\|\left(\boldsymbol{u}_{i} \boldsymbol{S} \boldsymbol{B}^{T}-\boldsymbol{a}_{i}\right) \boldsymbol{W}_{i}\right\|_{2}^{2}+\frac{1}{\gamma}\left\|\boldsymbol{u}_{i}\right\|_{2}^{2}\right)\right) . \tag{5}
\end{equation*}
$$

The inner optimization problem $\min _{\boldsymbol{u}_{i}}\left\|\left(\boldsymbol{u}_{i} \boldsymbol{S} \boldsymbol{B}^{T}-\boldsymbol{a}_{i}\right) \boldsymbol{W}_{i}\right\|_{2}^{2}+\frac{1}{\gamma}\left\|\boldsymbol{u}_{i}\right\|_{2}^{2}$ can be solved in closed form given $\boldsymbol{S}$, as it is a weighted linear regression problem with Tiknorov regularization, see Bertsimas and van Parys (2020). The closed form solution is:

$$
\begin{equation*}
\min _{\boldsymbol{u}_{i}}\left\|\left(\boldsymbol{u}_{i} \boldsymbol{S} \boldsymbol{B}^{T}-\boldsymbol{a}_{i}\right) \boldsymbol{W}_{i}\right\|_{2}^{2}+\frac{1}{\gamma}\left\|\boldsymbol{u}_{i}\right\|_{2}^{2}=\overline{\boldsymbol{a}}_{i}\left(\boldsymbol{I}_{m}+\gamma \boldsymbol{W}_{i} \boldsymbol{B} \boldsymbol{S} \boldsymbol{B}^{T} \boldsymbol{W}_{i}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T} . \tag{6}
\end{equation*}
$$

So Problem (5) can be simplified to:

$$
\min _{\boldsymbol{s} \in S_{k}^{p}} \frac{1}{n m}\left(\sum_{i=1}^{n} \overline{\boldsymbol{a}}_{i}\left(\boldsymbol{I}_{m}+\gamma \boldsymbol{W}_{i} \boldsymbol{B} \boldsymbol{S} \boldsymbol{B}^{T} \boldsymbol{W}_{i}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}\right) .
$$

Finally, we notice that

$$
\boldsymbol{B} \boldsymbol{S} \boldsymbol{B}^{T}=\sum_{j=1}^{p} s_{j} \boldsymbol{b}^{j}\left(\boldsymbol{b}^{j}\right)^{T}=\sum_{j=1}^{p} s_{j} \boldsymbol{K}_{j},
$$

and we obtain the required expression. Since $K_{j}$ are positive semi-definite, and the inverse of positive semi-definite matrices is a convex function, the entire function is convex in $s$.

With Theorem 1, our original problem can now be restated as:

$$
\begin{equation*}
\min _{\boldsymbol{s} \in S_{k}^{p}} c(\boldsymbol{s})=\frac{1}{n m} \sum_{i=1}^{n} \overline{\boldsymbol{a}}_{i}\left(\boldsymbol{I}_{m}+\gamma \boldsymbol{W}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{K}_{j}\right) \boldsymbol{W}_{i}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T} . \tag{7}
\end{equation*}
$$

This can be solved utilizing the cutting plane algorithm first introduced by Duran and Grossmann (1986), summarized as Algorithm 1 .

```
Algorithm 1 Cutting-plane algorithm for matrix completion with side information.
    procedure CUTPLANES \((\boldsymbol{A}, \boldsymbol{B}) \quad\) \# masked matrix \(\boldsymbol{A}\), and feature matrix \(\boldsymbol{B}\)
        \(t \leftarrow 1\)
        \(s_{1} \leftarrow\) warm start \# Heuristic Warm Start
        \(\eta \leftarrow 0 \quad\) \# Initialize feasible solution variable
        while \(\eta_{t}<c\left(\boldsymbol{s}_{t}\right)\) do \# While the current solution is not optimal
            \(\boldsymbol{s}_{t+1}, \eta_{t+1} \leftarrow \underset{\boldsymbol{s} \in S_{k}^{p}, \eta>0}{\arg \min } \quad \eta\) s.t. \(\quad \eta \geq c\left(\boldsymbol{s}_{i}\right)+\nabla c\left(\boldsymbol{s}_{i}\right)^{T}\left(\boldsymbol{s}-\boldsymbol{s}_{i}\right) \quad \forall i \in[t]\)
            \(t \leftarrow t+1\)
        end while
        \(s \leftarrow s_{t}\)
        \(i \leftarrow 1\)
        for \(i<n\) do \# Fill each row \(\boldsymbol{x}_{i}\) of final output matrix \(\boldsymbol{X}\)
            \(\boldsymbol{x}_{i} \leftarrow \boldsymbol{B}^{\boldsymbol{s}}\left(\left(\boldsymbol{B}^{\boldsymbol{s}}\right)^{T} \boldsymbol{W}_{i} \boldsymbol{B}^{\boldsymbol{s}}\right)^{-1}\left(\boldsymbol{B}^{\boldsymbol{s}}\right)^{T} \overline{\boldsymbol{a}}_{i}^{T} \# \boldsymbol{B}^{\boldsymbol{s}}\) is submatrix of \(\boldsymbol{B}\) with \(\boldsymbol{s}\) columns
        end for
        return \(\boldsymbol{X}\) \# Return the filled matrix \(\boldsymbol{X}\)
    end procedure
```

The cutting plane algorithm, at iteration $t$, adds a linear approximation of $c(s)$ at the current feasible solution $s_{t}$ to the set of constraints:

$$
\begin{equation*}
\eta \geq c\left(\boldsymbol{s}_{t}\right)+\nabla c\left(\boldsymbol{s}_{t}\right)^{T}\left(\boldsymbol{s}-\boldsymbol{s}_{t}\right) \tag{8}
\end{equation*}
$$

and we solve the mixed-integer linear programming problem:

$$
\begin{aligned}
& \min _{s \in S_{k}^{p}, \eta \geq 0} \eta \\
& \eta \geq c\left(\boldsymbol{s}_{t}\right)+\nabla c\left(\boldsymbol{s}_{t}\right)^{T}\left(\boldsymbol{s}-\boldsymbol{s}_{t}\right), \quad i \in[t]
\end{aligned}
$$

to obtain $s_{t+1}, \eta_{t+1}$. We see that $\eta_{t+1}$ is exactly the minimum value of the current approximation of $c(\boldsymbol{s}), c_{t}(\boldsymbol{s})$, defined below:

$$
\eta_{t+1}=\min _{s} \max _{i \in[t]} c\left(\boldsymbol{s}_{t}\right)+\nabla x\left(\boldsymbol{s}_{t}\right)^{T}\left(\boldsymbol{s}-\boldsymbol{s}_{t}\right)=\min _{s} c_{t}(\boldsymbol{s}) .
$$

Since $c(\boldsymbol{s})$ is convex, the piecewise linear approximation $c_{t}(\boldsymbol{s})$ is an outer approximation $\left(c_{t}(\boldsymbol{s})<=\right.$ $c(s) \forall s)$, so $\eta_{t} \leq c\left(s_{t}\right) \forall t$. As the algorithm progresses, the set of linear approximations form an increasingly better approximation of $c(s)$, and $\eta_{t}$ increases with $t$. The algorithm terminates once $\eta_{t}$ does not further increase, as it implies the linear approximation shares the same minimum value as the true function $c(s)$, which is the desired value.

Once the optimal solution $s^{*}$ is reached, we can obtain the optimal $\boldsymbol{U}$ using the closed form solution in (6) and recover $\boldsymbol{X}$. In the next section, we discuss how this algorithm can be implemented in the context of $c(s)$ in (7) and derive its computational complexity.

### 2.2. Implementation and Computational Complexity of CutPlanes

The computational complexity of the cutting plane comes from calculating $c(\boldsymbol{s})$ and its derivative $\nabla c(s)$. We first introduce the notations $\alpha_{i}(s) \in \mathbb{R}$ and $\gamma_{i}(s) \in \mathbb{R}^{m \times 1}$.

$$
\begin{equation*}
\alpha_{i}(\boldsymbol{s})=\frac{1}{m} \overline{\boldsymbol{a}}_{i} \gamma_{i}(\boldsymbol{s})=\frac{1}{m} \overline{\boldsymbol{a}}_{i}\left[\left(\boldsymbol{I}_{m}+\gamma \boldsymbol{W}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{K}_{j}\right) \boldsymbol{W}_{i}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}\right], \quad i=1, \ldots, n . \tag{9}
\end{equation*}
$$

Then, the function $c(s)$ in (7) can be expressed as

$$
\begin{equation*}
c(\boldsymbol{s})=\frac{1}{n} \sum_{i=1}^{n} \alpha_{i}(\boldsymbol{s})=\frac{1}{n} \sum_{i=1}^{n} \frac{\overline{\boldsymbol{a}}_{i} \gamma_{i}(\boldsymbol{s})}{m} . \tag{10}
\end{equation*}
$$

To calculate the derivative $\nabla c(s)$, it is easier to utilize the expression in Theorem 1 and then utilize the chain rule. After some manipulations, we obtain

$$
\begin{equation*}
\nabla c(\boldsymbol{s})=\frac{1}{n} \sum_{i=1}^{n}-\frac{\gamma\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}(\boldsymbol{s})\right)^{2}}{m} . \tag{11}
\end{equation*}
$$

Therefore, we would focus on calculating $\gamma_{i}(s)$. First, by the Matrix Inversion Lemma (Woodbury (1949)) we have

$$
\begin{align*}
\boldsymbol{\gamma}_{i}(\boldsymbol{s}) & =\left(\boldsymbol{I}_{m}+\gamma \boldsymbol{W}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{K}_{j}\right) \boldsymbol{W}_{i}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T} \\
& =\left(\boldsymbol{I}_{m}-\boldsymbol{V}\left(\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}^{T} \boldsymbol{W}_{i} \boldsymbol{V}\right)^{-1} \boldsymbol{V}^{T}\right) \overline{\boldsymbol{a}}_{i}^{T} \\
& =\left(\overline{\boldsymbol{a}}_{i}^{T}-\boldsymbol{V}\left(\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}^{T} \boldsymbol{W}_{i} \boldsymbol{V}\right)^{-1} \boldsymbol{V}^{T} \overline{\boldsymbol{a}}_{i}^{T}\right), \tag{12}
\end{align*}
$$

where $\boldsymbol{V} \in \mathbb{R}^{m \times k}$ is the feature matrix formed by the $k$ columns of $\boldsymbol{B}$ such that $s_{j}=1$, and we have suppressed the dependency of $\boldsymbol{V}$ on $\boldsymbol{s}$ for notation ease. Note that in order to compute $\gamma_{i}(\boldsymbol{s})$ using Eq. (91) we need to invert an $m \times m$ matrix, while from Eq. (12) we need to invert a $k \times k$ matrix $\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}^{T} \boldsymbol{W}_{i} \boldsymbol{V}$, which only requires $O\left(k^{3}\right)$ calculations. Furthermore, note that calculating $\overline{\boldsymbol{a}}_{i} \overline{\boldsymbol{a}}_{i}^{T}$ only requires $\left|\Omega_{i}\right|$ multiplications where $\Omega_{i}$ is the number of known entries in row $i$ of $\boldsymbol{A}$ as we do not need to multiply on the unknown entries. Similarly, we can compute $\boldsymbol{V}^{T} \overline{\boldsymbol{a}}_{i}^{T}$ in $\left|\Omega_{i}\right| k$ multiplications, and $\boldsymbol{V}^{T} \boldsymbol{W}_{i} \boldsymbol{V}$ in $\left|\Omega_{i}\right| k^{2}$ multiplications.

Therefore, we can compute $\gamma_{i}(s)$ in floating point complexity of $O\left(\left|\Omega_{i}\right| k^{2}+k^{3}\right)$. Then to calculate $\overline{\boldsymbol{a}}_{i} \boldsymbol{\gamma}_{i}(\boldsymbol{s})$ in (10) and $-\gamma\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \gamma_{i}(\boldsymbol{s})\right)^{2}$ (11) only requires $O\left(\left|\Omega_{i}\right|\right)$ and $O\left(\left|\Omega_{i}\right| p\right)$ calculations respectively. Thus, the total complexity of generating a full cutting plane is:

$$
\begin{equation*}
\sum_{i=1}^{n} O\left(\left|\Omega_{i}\right| p+\left|\Omega_{i}\right| k^{2}+k^{3}\right)=O\left(|\Omega|\left(p+k^{2}\right)+n k^{3}\right) . \tag{13}
\end{equation*}
$$

### 2.3. Two-sided Information Case

In this section, we briefly discuss the matrix completion problem under the two-sided information case, and how it reduces to the problem of sparse linear regression. The two sided interpretable matrix completion problem with Tikhonov regularization can be stated as follows:

$$
\begin{equation*}
\min _{\boldsymbol{L}} \frac{1}{n m}\left(\sum_{(i, j) \in \Omega}\left(X_{i j}-A_{i j}\right)^{2}+\frac{1}{\gamma}\|\boldsymbol{L}\|_{2}^{2}\right) \quad \text { subject to } \quad \boldsymbol{X}=\boldsymbol{U} \boldsymbol{L} \boldsymbol{B}^{T} \quad\|\boldsymbol{L}\|_{0}=k, \tag{14}
\end{equation*}
$$

where $\boldsymbol{U} \in \mathbb{R}^{n \times p_{1}}$ is a known matrix of $p_{1}$ features of each row, $\boldsymbol{B} \in \mathbb{R}^{m \times p_{2}}$ is a known matrix of $p_{2}$ features of each column, and $\boldsymbol{L} \in \mathbb{R}^{p_{1} \times p_{2}}$ is a sparse matrix that has $k$ nonzero entries, ensuring that $\operatorname{Rank}(\boldsymbol{X}) \leq k$. We note that in Eq. (14) we restrict the support of matrix $\boldsymbol{L}$ to be $k$, rather than forcing the entries of $\boldsymbol{L}$ to be binary. This is because unlike in the one-sided case, both $\boldsymbol{U}$ and $\boldsymbol{B}$ are known, so we cannot apply the scaling transformation in (3).

We denote by $\boldsymbol{u}^{i} \in \mathbb{R}^{n \times 1}$ the $i$ th column of $\boldsymbol{U}$ and $\boldsymbol{b}^{j} \in \mathbb{R}^{m \times 1}$ the $j$ th column of $\boldsymbol{B}$. We introduce the matrices $\boldsymbol{W}_{i}$ as in Theorem 1 Using $\boldsymbol{X}=\boldsymbol{U} \boldsymbol{L} \boldsymbol{B}$, we can write

$$
X_{i j}=\sum_{q=1}^{p_{1}} \sum_{\ell=1}^{p_{2}} L_{q, \ell} D_{i j}^{q, \ell}
$$

where $D_{i j}^{q, \ell}=\left(\boldsymbol{u}^{q}\left(\boldsymbol{b}^{\ell}\right)^{T}\right)_{i j}$ is the $(i, j)$ th entry of the matrix formed by multiplying $q$ th column of $\boldsymbol{U}$ with $\ell$ th column of $\boldsymbol{B}$. Then, Problem (14) becomes:

$$
\begin{equation*}
\min _{L} \frac{1}{n m}\left(\sum_{(i, j) \in \Omega}\left(\sum_{q=1}^{p_{1}} \sum_{\ell=1}^{p_{2}} L_{q, \ell} D_{i j}^{q, \ell}-A_{i j}\right)^{2}+\frac{1}{\gamma}\|\boldsymbol{L}\|_{2}^{2}\right) \quad \text { subject to } \quad\|\boldsymbol{L}\|_{0}=k . \tag{15}
\end{equation*}
$$

As every $\boldsymbol{D}$ matrix is known, this becomes a sparse regression problem where there are $p_{1} p_{2}$ features to choose from (the $\boldsymbol{D}$ matrices), there are $|\Omega|$ samples (the $\boldsymbol{A}$ matrix), the sparsity requirement is $k$, the regression coefficients are $\boldsymbol{L}$, and we have Tikhonov regularization. Vectorizing $\boldsymbol{D}, \boldsymbol{L}$, and $\boldsymbol{A}$ reduces the problem back to the familiar form of sparse linear regression, that can be solved by the algorithm developed in Bertsimas and van Parys (2020) at scale.

## 3. OptComplete: The Stochastic Cutting Plane Speedup

In this section, we introduce OptComplete, a stochastic version of the cutting plane algorithm introduced in Section2 We present theoretical results to show that the stochastic algorithm recovers the true optimal solution of the original algorithm with high probability without distributional assumptions. We also include a discussion on the dependence of such probability with various factors and its favorable theoretical computational complexity.

### 3.1. Introduction of OptComplete

In the previous section, we showed that through careful evaluation, we can calculate a full cutting plane in $O\left(|\Omega|\left(p+k^{2}\right)+n k^{3}\right)$ calculations. However, in very high dimensions where $|\Omega|, n, m$ are extremely large, the cost of generating the full cutting plane is still prohibitive. Thus, we consider generating approximations of the cutting plane that would enable the algorithm to scale for high values for $n$ and $m$. Specifically, consider the cutting plane function in (10), reproduced below:

$$
c(s)=\frac{1}{n} \sum_{i=1}^{n} \alpha_{i}(s)
$$

where:

$$
\alpha_{i}(s)=\frac{1}{m}\left(\overline{\boldsymbol{a}}_{i} \overline{\boldsymbol{a}}_{i}^{T}-\overline{\boldsymbol{a}}_{i} \boldsymbol{V}\left(\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}^{T} \boldsymbol{W}_{i} \boldsymbol{V}\right)^{-1} \boldsymbol{V}^{T} \overline{\boldsymbol{a}}_{i}^{T}\right) .
$$

We approximate the inner term $\alpha_{i}(\boldsymbol{s})$ by choosing $1 \leq f<m$ samples from $\{1, \ldots, m\}$ without replacement, with the set denoted $F$. Then we formulate the submatrix $\boldsymbol{V}_{F}$ with such selected rows, and similarly with $\overline{\boldsymbol{a}}_{F i}$. Then we calculate the approximation:

$$
\alpha_{i}(\boldsymbol{s}) \approx \alpha_{i}^{F}(\boldsymbol{s})=\frac{1}{f}\left(\overline{\boldsymbol{a}}_{F i} \overline{\boldsymbol{a}}_{F i}^{T}-\overline{\boldsymbol{a}}_{F i} \boldsymbol{V}_{F}\left(\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}_{F}^{T} \boldsymbol{W}_{i} \boldsymbol{V}_{F}\right)^{-1} \boldsymbol{V}_{F}^{T} \overline{\boldsymbol{a}}_{F i}^{T}\right) .
$$

Then we choose $1 \leq g<n$ samples from $\{1, \ldots, n\}$ without replacement, with the set denoted $G$.
We can then calculate an approximation of $c(s)$ using the approximated $\alpha_{i}^{F}(s)$ :

$$
c(s) \approx \tilde{c}_{G}^{F}(s)=\frac{1}{r} \sum_{i \in G} \alpha_{i}^{F}(s),
$$

where the set $F$ is chosen independently for every row $i \in G$. Then the derivative of $\tilde{c}_{G}^{F}(s)$ is:

$$
\nabla \tilde{c}_{G}^{F}(s)=\frac{1}{r} \sum_{i \in G} \frac{\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}^{F}(s)\right)^{2}}{f}
$$

Using such approximations, we can derive a stochastic cutting-plane algorithm, which we call OptComplete presented as Algorithm 2.

For this algorithm to work, we need the approximation $\tilde{c}_{G}^{F}(\boldsymbol{s})$ and its derivative to be close to the nominal values. Furthermore, the approximated cutting planes should not cutoff the true solution. In the next section, we show that OptComplete enjoys such properties with high probability. In Section 3.3, we discuss how to select the size of $f$ and $g$.

### 3.2. Main Theoretical Results

We would first show that the inner approximation is close to the true term with high probability:

Theorem 2 Let $\boldsymbol{A}$ be a partially known matrix, $\boldsymbol{B}$ a known feature matrix, and $\boldsymbol{W}_{i}$ as defined in Theorem 团, Let $F$ be a random sample of size $f$ from the set $\{1, \ldots, m\}$, chosen without replacement. With probability at least $1-\epsilon$, we have

$$
\begin{array}{r}
\left|\alpha_{i}(s)-\alpha_{i}^{F}(s)\right| \leq \sqrt{\frac{M k \log \left(\frac{k}{\epsilon}\right)}{f}}, \quad \forall i \in\{1, \ldots, m\}, \quad \forall s \in S_{k}^{p}, \\
\left\|\frac{\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}(\boldsymbol{s})\right)^{2}}{m}-\frac{\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}^{F}(\boldsymbol{s})\right)^{2}}{f}\right\|_{2} \leq \sqrt{\frac{M^{\prime}(p+k) \log \left(\frac{k}{\epsilon}\right)}{f}}, \quad \forall i \in\{1, \ldots, m\}, \quad \forall s \in S_{k}^{p},
\end{array}
$$

where $M, M^{\prime}$ are absolute constants.

```
Algorithm 2 Stochastic Cutting-plane algorithm for matrix completion with side information.
    procedure OptComplete \((\boldsymbol{A}, \boldsymbol{B}) \quad \#\) masked matrix \(\boldsymbol{A}\), and feature matrix \(\boldsymbol{B}\)
    \(t \leftarrow 1\)
    \(s_{1} \leftarrow\) random initialization
    \(\eta \leftarrow 0 \quad\) \# Initialize feasible solution variable
    while \(\eta_{t}<c\left(\boldsymbol{s}_{t}\right)\) do \# While the current solution is not optimal
            \(G \leftarrow|g|\)-sized random sample of \(\{1, \ldots, n\}\) with replacement
            for \(i \in G\) do \# Generate \(F\) for each row in random sample
                \(F_{i} \leftarrow|f|\)-sized random sample of \(\{1, \ldots, m\}\) with replacement
            end for
            \(\boldsymbol{s}_{t+1}, \eta_{t+1} \leftarrow \underset{s \in S_{k}^{P}, \eta>0}{\arg \min } \quad \eta\) s.t. \(\quad \eta \geq \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{i}\right)+\nabla \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{i}\right)^{T}\left(\boldsymbol{s}-\boldsymbol{s}_{i}\right) \quad \forall i \in[t]\)
            \(t \leftarrow t+1\)
    end while
    \(s \leftarrow s_{t}\)
    \(i \leftarrow 1\)
    for \(i<n\) do \# Fill each row \(\boldsymbol{x}_{i}\) of final output matrix \(\boldsymbol{X}\)
        \(\boldsymbol{x}_{i} \leftarrow \boldsymbol{B}^{\boldsymbol{s}}\left(\left(\boldsymbol{B}^{\boldsymbol{s}}\right)^{T} \boldsymbol{W}_{i} \boldsymbol{B}^{\boldsymbol{s}}\right)^{-1}\left(\boldsymbol{B}^{\boldsymbol{s}}\right)^{T} \overline{\boldsymbol{a}}_{i}^{T} \# \boldsymbol{B}^{\boldsymbol{s}}\) is submatrix of \(\boldsymbol{B}\) with \(\boldsymbol{s}\) columns
        end for
        return \(\boldsymbol{X}\) \# Return the filled matrix \(\boldsymbol{X}\)
    end procedure
```

We see that, without assumptions on the data, the inner approximation for both the value and the derivative follows a bound with $O\left(\sqrt{\frac{(p+k)}{f}}\right)$ terms with very high probability. Furthermore, inverting the statements give that, for all $i \in\{1, \ldots, m\}$ and all $s \in S_{k}^{p}$ :

$$
\begin{gathered}
\mathbb{P}\left(\left|\alpha_{i}(s)-\alpha_{i}^{F}(s)\right| \geq \delta\right) \leq k \exp \left(-\frac{f \delta^{2}}{M k}\right), \\
\mathbb{P}\left(\left\|\frac{\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}(\boldsymbol{s})\right)^{2}}{m}-\frac{\left(\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}^{F}(\boldsymbol{s})\right)^{2}}{f}\right\|_{2} \geq \delta\right) \leq k \exp \left(-\frac{f \delta^{2}}{M^{\prime}(p+k)}\right) .
\end{gathered}
$$

So the failure probability drops off exponentially with increasing bound $\delta$, reflecting a Gaussian tail structure for $\alpha_{i}(s)$ and $\boldsymbol{B}^{T} \boldsymbol{W}_{i} \boldsymbol{\gamma}_{i}^{F}(s)$. The proof is contained in Appendix A

Using this result, we are able to prove a tight deviation bound for the approximated cost function $c_{G}^{F}(\boldsymbol{s})$ and $\nabla c_{G}^{F}(\boldsymbol{s})$ :

Theorem 3 Let $\boldsymbol{A}$ be a partially known matrix, $\boldsymbol{B}$ a known feature matrix, and $\boldsymbol{W}_{i}$ as defined in Theorem 1. Let $G$ be a random sample of size $g$ from $\{1, \ldots, n\}$ chosen without replacement. Then for each $i \in G$, we let $F_{i}$ be a random sample of size $f$ from the set $\{1, \ldots, m\}$, all chosen without replacement. We have, with probability at least $1-\epsilon$ :

$$
\begin{gathered}
\left|\tilde{c}_{G}^{F}(s)-c(s)\right| \leq \sqrt{\frac{A k \log \left(\frac{k}{\epsilon}\right)}{g}}, \quad \forall s \in S_{k}^{p}, \\
\left\|\nabla \tilde{c}_{G}^{F}(s)-\nabla c(s)\right\|_{2} \leq \sqrt{\frac{B(p+k) \log \left(\frac{k}{\epsilon}\right)}{g}}, \quad \forall s \in S_{k}^{p},
\end{gathered}
$$

where $A, B$ are absolute constants.

Similar to the inner approximations, $c_{G}^{F}(\boldsymbol{s})$ and $\nabla c_{G}^{F}(\boldsymbol{s})$ has Gaussian tails. Furthermore, the scaling here only depends on $g$ and not $f$ : This shows that the error of the inner approximation is dominated by the outer sampling of the rows $G$. The proof is contained in Appendix B

Then, using this result, we are able to prove our main result for OptComplete. We would first introduce a new definition:

Definition 1 The convexity parameter a of the cost function $c(s)$ is defined as the largest positive number for which the the following statement is true:

$$
\begin{equation*}
c(\boldsymbol{s}) \geq c\left(\boldsymbol{s}_{0}\right)+\nabla c\left(\boldsymbol{s}_{0}\right)^{T}\left(s-\boldsymbol{s}_{0}\right)+\frac{a^{2}}{2}\left(s-s_{0}\right)^{T}\left(\boldsymbol{s}-\boldsymbol{s}_{0}\right) \quad \forall \boldsymbol{s}, \boldsymbol{s}_{0} \in S_{k}^{p} \forall i \tag{16}
\end{equation*}
$$

We have the following proposition which shows that unless the cost function is degenerate (i.e. different sets of $k$ features doesn't change the solution), we always have a positive convexity parameter:

Proposition 1 Assume that there does not exist $\boldsymbol{s}_{1}, \boldsymbol{s}_{0} \in S_{k}^{p}$ such that $c\left(\boldsymbol{s}_{1}\right)=c\left(\boldsymbol{s}_{0}\right)$. Then $a>0$.

The proof is contained in Appendix C. Now, we state our main theorem for OptComplete.

Theorem 4 For the matrix completion problem (4), let $\boldsymbol{B} \in \mathbb{R}^{m \times p}$ be a known feature matrix, $\boldsymbol{A} \in \mathbb{R}^{n \times m}$ a matrix with entries partially known, and OptComplete as defined in Algorithm 园, Assume that Problem (4) is feasible. Then, OptComplete terminates in a finite number of steps $C$, and finds an optimal solution of (4) with probability at least $1-k C \exp \left(-\frac{D a^{4} g}{(p+k)}\right)$ where $D$ is an absolute constant independent of $C, f, g, k, p$, and $a$ is the convexity parameter of the functions $\tilde{\alpha}_{i}^{s}(s)$.

The proof is contained in Appendix D. This theorem shows that as long as the original problem is feasible, OptComplete is able to find the optimal solution of the original binary convex problem with exponentially vanishing failure probability that scales as $O\left(\exp \left(\frac{-g}{(p+k)}\right)\right)$. The theorem requires no assumptions on the data, and thus applies generally. We again note that the bound does not depend on $f$ and only on $g$ : we would discuss how this would inform our selection of the size of $f$ and $g$ in the next section.

### 3.3. Sampling Size and Computational Complexity

To select an appropriate $f$ and $g$, we first note that Candès and Tao (2010) showed that to complete a square $N \times N$ matrix of rank $k$, we need at least $O(k N \log N)$ elements. Assume an average known rate of $\alpha=\frac{|\Omega|}{m n}$ in the original matrix $\boldsymbol{A}$, the expected number of known elements under a sampling of $f$ and $g$ is $\alpha f g$. Using $N^{2}=m n$, we need that:

$$
\begin{equation*}
\alpha f g \geq c \cdot k \sqrt{n m} \log (\sqrt{n m}) \tag{17}
\end{equation*}
$$

for some constant $c$. Theorem [4 showed that the bound on failure probability scales with $O\left(\exp \left(\frac{-g}{(p+k)}\right)\right)$, and thus we cannot have $g$ too small. Using (13), the complexity of the cutting plane with $f$ and $g$ samples are:

$$
\begin{equation*}
O\left(\alpha f g\left(p+k^{2}\right)+g k^{3}\right) . \tag{18}
\end{equation*}
$$

Therefore, if we fix the expected known elements ( $\alpha f g$ ) constant, it is more advantageous to select a smaller $g$, as $g$ scales with $k^{3}$. Thus, we set:

$$
\begin{equation*}
f=\min \left(\frac{c k \sqrt{m n} \log (\sqrt{m n})}{\alpha \min \left(g_{0}, n\right)}, m\right), \quad g=\min \left(g_{0}, n\right) \tag{19}
\end{equation*}
$$

Experimentally, we found $g_{0}=100, c=1$ to generate good results (the results were similar for $\frac{1}{2} \leq c \leq 2$ ). Therefore, by (18), the approximated cutting plane has a computational complexity of:

$$
\begin{equation*}
O\left(k \sqrt{m n} \log (\sqrt{m n})\left(p+k^{2}\right)\right) . \tag{20}
\end{equation*}
$$

This scales in a square root fashion in $n$ and $m$, rather than linearly in $n$ and $m$ for the full cutting plane. This allows OptComplete to enjoy a considerable speedup compared to CutPlanes, as demonstrated in Section 4 .

## 4. Synthetic Data Experiments

We assume that the matrix $\boldsymbol{A}=\boldsymbol{U} \boldsymbol{V}+\boldsymbol{E}$, where $\boldsymbol{U} \in \mathbb{R}^{n \times k}, \boldsymbol{V} \in \mathbb{R}^{k \times m}$, and $\boldsymbol{E}$ is an error matrix with individual elements sampled from $N(0,0.01)$. We sample the elements of $\boldsymbol{U}$ and $\boldsymbol{V}$ from a uniform distribution of $[0,1]$, and then randomly select a fraction $\mu=1-\alpha$ to be missing. We formulate the feature matrix $\boldsymbol{B}$ by combining $\boldsymbol{V} \in \mathbb{R}^{k \times m}$ with a confounding matrix $\boldsymbol{Z} \in$ $\mathbb{R}^{(p-k) \times m}$ that contains unnecessary factors sampled similarly from the Uniform $[0,1]$ distribution. We run OptComplete on a server with 16 CPU cores, using Gurobi 8.1.0. For each combination ( $m, n, p, k, \mu$ ), we ran 10 tests and report the median value for every statistic.

We report the following statistics with $s^{*}$ being the ground-truth factor vector, and $\bar{s}$ the estimated factor vector.

- $n, m$ - the dimensions of $\boldsymbol{A}$.
- $p$ - the number of features in the feature matrix.
- $k$ - the true number of features.
- $\mu$ - The fraction of missing entries in $\boldsymbol{A}$.
- $T$ - the total time taken for the algorithm.
- MAPE - the Mean Absolute Percentage Error (MAPE) for the retrieved matrix $\hat{\boldsymbol{A}}$ :

$$
\mathrm{MAPE}=\frac{1}{|\mathcal{S}|} \sum_{(i, j) \in \mathcal{S}} \frac{\left|\hat{A}_{i j}-A_{i j}\right|}{\left|A_{i j}\right|}
$$

where $\mathcal{S}=\Omega^{c}$ is the set of missing data in $\boldsymbol{A}$.
Since the concept of Interpretable Matrix Completion is new, there is a lack of directly comparable algorithms in the literature. Thus, in lieu, we compare OptComplete to state-of-the-art solvers for Inductive Matrix Completion and general matrix completion, which are:

- IMC by Natarajan and Dhillon (2014) - This algorithm is a well-accepted benchmark for testing Inductive Matrix Completion algorithms.
- SoftImpute-ALS (SIALS) by Hastie et al. (2015) - This is widely recognized as a state-of-theart matrix completion method without feature information. It has among the best scaling behavior across all classes of matrix completion algorithms as it utilizes fast alternating least squares to achieve scalability.

We use the best existing implementations of IMC (Matlab 2018b) and SIALS (R 3.4.4, package softImpute) with parallelization on the same server.

We further compare our algorithm to CutPlanes, the original cutting plane algorithm developed in Section2. It is known that for general mixed-integer convex problems, the cutting plane algorithm has the best overall performance (see e.g. Lubin et al. (2016) for details), and thus CutPlanes represent a good baseline of comparison for OptComplete.

We randomly selected $20 \%$ of those elements masked to serve as a validation set. The regularization parameter $\gamma$ of OptComplete, the rank parameter of IMC and the penalization parameter $\lambda$ of IMC and SIALS are selected using the validation set. The results are separated into sections below. The first five sections modify one single variable out of $n, m, p, k, \mu$ to investigate OptComplete's scalability, where the leftmost column indicates the variable modified. The last section compares the four algorithms scalability for a variety of parameters that reflect more realistic scenarios.


Table 1 Comparison of OptComplete, IMC and SIALS on synthetic data. $N / A$ means the algorithm did not complete running in 20 hours, corresponding to 72000 seconds.

Overall, we see that OptComplete achieves near-exact retrieval on all datasets evaluated, and successfully recovers the factors in the ground truth. The solutions (and its error) also matches with that of CutPlanes, the standard cutting plane algorithm. The non-zero MAPE is due to the random noise added resulting in slightly perturbed coefficients.

For the realistic and large data sizes in the last panel, we see that OptComplete not only achieves near-exact retrieval, it does so while requiring considerably less time than IMC and SIALS at the same time. For $m, n$ on the scale of $n=10^{6}$ and $m=10^{4}$, OptComplete is over 20 times faster than IMC and over 40 times faster than SIALS. At the scale of $n=10^{6}$ and $m=10^{5}$, IMC and SIALS did not finish running within 20 hours, while OptComplete completed in just under 12 minutes. We also see that OptComplete achieves very significant speedups compared to the standard cutting plane algorithm - up to 60 x at the scale of $n=10^{5}$ and $m=10^{4}$.

We analyze the scaling of OptComplete as a function of:

1. $\mu$ - The algorithm is able to retrieve the exact factors used even with $95 \%$ of missing data. Furthermore, the running time decreased with increasing missing entries, consistent with the fact that is computational complexity scales with $|\Omega|$.
2. $n$ - The algorithm has good scalability in $n$, reflecting its $O(\sqrt{n} \log (n))$ type complexity. This allows the algorithm to support matrices with $n$ in the $10^{6}$ range. Its scaling behavior is superior to both IMC and SIALS.
3. $m$ - The algorithm scales exceptionally well in $m$. We observe that empirically the algorithm runtime seems to grow much slower than the theoretical $O(\sqrt{m} \log (m))$ dependence. A closer examination reveals that as $m$ increases, the number of cutting planes generated by Gurobi is decreasing. Qualitatively, this can be explained by a larger $m$ giving the algorithm more signal to find which $k$ features are the correct ones out of the $p$ ones. We note that such behavior is also exhibited by CutPlanes, as it roughly scales as $O(\sqrt{m})$ rather than the $O(m)$ as expected. We see that IMC and SIALS scales as $O(m)$.
4. $p$ - The algorithm scales relatively well in $p$, which reflects the performance of the Gurobi solver. We empirically observe that Gurobi is generating roughly $O(1)-O(p)$ cutting planes. Thus, as each cutting plane is $O(p)$, we expect $O(p)-O\left(p^{2}\right)$ dependence, as is observed here. We note that OptComplete achieves similar scaling behavior as IMC in $p$. Note here the SIALS algorithm does not utilize feature information and thus a change in $p$ does not affect the algorithm's run speed.
5. $k$ - The algorithm does not scale very well in $k$. We empirically observe that Gurobi solver is roughly generating $O(k)$ cutting planes and each cutting plane has cubic dependence on $k$. It appears that SIALS and IMC almost have a linear scaling behavior. However, in most applications, such as recommendation systems or low-rank retrieval, $k$ is usually kept very low ( $k \leq 30$ ), so this is not a particular concern.

## 5. Real-World Experiments

In this section, we report on the performance of OptComplete on the Netflix Prize dataset (Bennett et al. 2007). This dataset was released in a competition to predict ratings of customers on unseen movies, given over 10 million ratings scattered across 500,000 people and 16,000 movies. Thus, when presented in a matrix $\boldsymbol{A}$ where $A_{i j}$ represents the rating of individual $i$ on movie $j$, the goal is to complete the matrix $\boldsymbol{A}$ under a low-rank assumption.

The feature matrix $\boldsymbol{B}$ of OptComplete is constructed using data from the TMDB Database, and covers 59 features that measure geography, popularity, top actors/actresses, box office, runtime, genre and more. The full list of 59 features is contained in Appendix E.

For this experiment, we included movies where all 59 features are available, and people who had at least 5 ratings present. This gives a matrix of 471,268 people and 14,538 movies. The slight reduction of size from the original data is due to the lack of features for about 2,000 niche movies. To observe the scalability of OptComplete, we created five data sets:

1. Base - $\boldsymbol{A}_{1}$ has dimensions $3,923 \times 103$.
2. Small - $\boldsymbol{A}_{2}$ has dimensions $18,227 \times 323$.
3. Medium - $\boldsymbol{A}_{3}$ has dimensions $96,601 \times 788$.
4. Large - $\boldsymbol{A}_{4}$ has dimensions $471,268 \times 1760$.
5. Full - $\boldsymbol{A}$ has dimensions $471,268 \times 14,538$.

These sizes are constructed such that the total number of elements in $\boldsymbol{A}$ in the successive sizes are approximately different by approximately an order of magnitude.

For each individual matrix, we uniformly randomly withhold $20 \%$ of the ratings as a test set $\mathcal{S}$, and use the remaining $80 \%$ of ratings to impute a complete matrix $\hat{\boldsymbol{A}}$ - we perform cross-validation on the appropriate hyperparameters. Then, we report MAPE.

For comparison, we again use IMC and SIALS. We set the maximum rank of SIALS to be $k$ the rank optimized for in OptComplete. The results are listed below:

| $\boldsymbol{n}$ | $\boldsymbol{*} \boldsymbol{m}$ | $\boldsymbol{p}$ | $\boldsymbol{k}$ | $\boldsymbol{\mu} \%$ | OptComplete |  | IMC |  | SIALS |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  |  |  |  | MAPE | $\boldsymbol{T}$ | MAPE | $\boldsymbol{T}$ | MAPE |  |
| 3,923 | 103 | 59 | 5 | $92.6 \%$ | 6.0 s | $29.4 \%$ | 0.6 s | $34.2 \%$ | 0.3 s | $31.2 \%$ |
| 18,227 | 323 | 59 | 5 | $94.8 \%$ | 12.2 s | $21.8 \%$ | 5.2 s | $29.1 \%$ | 4.1 s | $24.1 \%$ |
| 96,601 | 788 | 59 | 5 | $94.2 \%$ | 25.5 s | $20.9 \%$ | 38.1 s | $28.7 \%$ | 30.4 s | $21.3 \%$ |
| 471,268 | 1,760 | 59 | 5 | $93.6 \%$ | 102 s | $18.8 \%$ | 460 s | $24.6 \%$ | 430 s | $19.8 \%$ |
| 471,268 | 14,538 | 59 | 5 | $94.1 \%$ | 170 s | $15.7 \%$ | 3921 s | $21.5 \%$ | 5300 s | $16.7 \%$ |

Table 2 Comparison of methods on Netflix data for $k=5$.

| $n$ | $m$ | $p$ | $k$ | $\boldsymbol{\mu} \%$ | OptComplete |  | IMC |  | SIALS |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  | T | MAPE | T | MAPE | T | MAPE |
| 3,923 | 103 | 59 | 10 | 92.6\% | 11.0s | 30.4\% | 1.4 s | $36.7 \%$ | 0.8s | 35.8\% |
| 18,227 | 323 | 59 | 10 | 94.8\% | 20.3s | 24.0\% | 12.5 s | 32.5\% | 7.0s | 28.9\% |
| 96,601 | 788 | 59 | 10 | 94.2\% | 45.9s | 22.3\% | 84.2 s | 29.6\% | 50.7 s | 22.8\% |
| 471,268 | 1,760 | 59 | 10 | 93.6\% | 260s | 20.7\% | 1022s | 24.8\% | 870s | 20.7\% |
| 471,268 | 14,538 | 59 | 10 | 94.1\% | 380s | 19.6\% | 8704s | 23.1\% | 10240s | 20.0\% |

Table 3 Comparison of methods on Netflix data for $k=10$.

We can see that OptComplete outperforms both IMC and SIALS in accuracy across the datasets under different $k$; furthermore in the two largest datasets OptComplete ran 10x to 20x faster than IMC and SIALS. Here we see that an increase from $k=5$ to $k=10$ actually decreased out-of-sample performance as additional factors are actually not very helpful in predictive customer tastes. The decline for OptComplete and IMC were especially higher due to the fact that the possible factors are fixed and thus an increase in the number of factors caused some non-predictive factors to be included.

For the $k=5$ case, OptComplete identified the following as the top factors that influences an individual's rating:

- IMDB Rating
- Genre: Drama
- Released within last 10 years
- Number of Top 100 Actors
- Produced in US

These factors provide an intuitive explanation of the individual ratings of each customer in terms of a small number of factors, while exceeding the high predictive accuracy of SIALS.

## 6. Conclusions

We have presented OptComplete, a scalable algorithm to retrieve a low-rank matrix in the presence of side information. Compared with state of the art algorithms for matrix completion, OptComplete exceeds current benchmarks on both scalability and accuracy and provides insight on the factors that affect the ratings.

## Appendices

## A. Proof of Theorem [2

We first note that since $S_{k}^{p}$ is a finite set, we only need to prove the result for a particular $s \in S_{k}^{p}$, and it would apply for all $s$. Therefore, we would assume $s$ is fixed below. For simplicity, we would only demonstrate the proof for $\alpha_{i}(\boldsymbol{s})=\overline{\boldsymbol{a}}_{i} \gamma_{i}(s)$, as the one for $\left(\boldsymbol{B} \boldsymbol{W}_{i} \gamma_{i}(s)\right)^{2}$ follows in the same exact fashion. Furthermore, since we are only focusing on one particular $i \in\{1, \ldots, n\}$, we would drop all $i$ subscripts below for ease of notation. The quantities of interest are therefore:

$$
\begin{aligned}
\alpha(s) & =\frac{1}{m}\left(\overline{\boldsymbol{a}}^{T}-\overline{\boldsymbol{a}} \boldsymbol{V}\left(\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}^{T} \boldsymbol{W} \boldsymbol{V}\right)^{-1} \boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}\right) \\
\alpha^{F}(\boldsymbol{s}) & =\frac{1}{f}\left(\overline{\boldsymbol{a}}_{F} \overline{\boldsymbol{a}}_{F}^{T}-\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}\left(\frac{\boldsymbol{I}_{k}}{\gamma}+\boldsymbol{V}_{F}^{T} \boldsymbol{W} \boldsymbol{V}_{F}\right)^{-1} \boldsymbol{V}_{F}^{T} \overline{\boldsymbol{a}}_{F}^{T}\right) .
\end{aligned}
$$

First, let $\overline{\boldsymbol{V}}=\boldsymbol{W} \boldsymbol{V}$, and let us consider a reduced QR factorization of $\overline{\boldsymbol{V}}=\boldsymbol{Q} \boldsymbol{R}$ where $\boldsymbol{Q} \in \mathbb{R}^{m \times k}$ has orthogonal columns such that $\boldsymbol{Q}^{T} \boldsymbol{Q}=m \cdot \boldsymbol{I}_{k}$, and $\boldsymbol{R}_{i} \in \mathbb{R}^{k \times k}$. Note such definition implies $\|\boldsymbol{R}\|=O(1)$. Then, we would rewrite the terms as follows:

$$
\begin{aligned}
\alpha(\boldsymbol{s}) & =\frac{\overline{\boldsymbol{a}}^{T}}{m}-\frac{\overline{\boldsymbol{a}} \boldsymbol{V}}{m}\left(\frac{\boldsymbol{I}_{k}}{m \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}\right)^{-1} \frac{\boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}}{m}, \\
\alpha^{F}(s) & =\frac{\overline{\boldsymbol{a}}_{F} \overline{\boldsymbol{a}}_{F}^{T}}{f}-\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}\left(\frac{\boldsymbol{I}_{k}}{f \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right)^{-1} \frac{\boldsymbol{V}_{F}^{T} \overline{\boldsymbol{a}}_{F}^{T}}{f} .
\end{aligned}
$$

We note that

$$
\begin{array}{rlrl}
\overline{\boldsymbol{a}}^{T} & =\sum_{i=1}^{m} \bar{a}_{i}^{2}, & \overline{\boldsymbol{a}} \boldsymbol{V} & =\sum_{i=1}^{m} \bar{a}_{i} \boldsymbol{v}_{i}, \\
\overline{\boldsymbol{a}}_{F} \overline{\boldsymbol{a}}_{F}^{T}=\sum_{i \in F} \bar{a}_{i}^{2}, & \overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F} & =\sum_{i \in F} \bar{a}_{i} \boldsymbol{v}_{i} .
\end{array}
$$

Therefore, if we treat $\bar{a}_{1}^{2}, \ldots \bar{a}_{m}^{2}$ as a finite population, then $\overline{\boldsymbol{a}}_{F} \overline{\boldsymbol{a}}_{F}^{T}$ is a random sample of $f$ points drawn without replacement from that set, and similarly for $\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}$. Therefore, we can then utilize Hoeffding's inequality to bound the deviation of these terms, as reproduced below:

Proposition 2 (Hoeffding's Inequality) Let $\mathcal{X}=\left(x_{1}, \ldots, x_{n}\right)$ be a finite population of $N$ points and $X_{1}, \ldots, X_{n}$ be a random sample drawn without replacement from $\boldsymbol{X}$. Let

$$
a=\min _{1 \leq i \leq n} x_{i} \quad \text { and } \quad b=\max _{1 \leq i \leq n} x_{i} .
$$

Then, for all $\epsilon>0$, we have

$$
\begin{equation*}
\mathbb{P}\left(\left|\frac{\sum_{i=1}^{n} X_{i}}{n}-\mu\right| \geq \epsilon\right) \leq 2 \exp \left(-\frac{2 n \epsilon^{2}}{(b-a)^{2}}\right) . \tag{A1}
\end{equation*}
$$

For a proof, see for example Boucheron et al. (2013). Then, applying Proposition 2 to $\overline{\boldsymbol{a}}_{f} \overline{\boldsymbol{a}}_{f}^{T}, \overline{\boldsymbol{a}}_{f} \boldsymbol{V}_{f}$, and inverting the inequality, we have

$$
\begin{align*}
\mathbb{P}\left(\left|\frac{\overline{\boldsymbol{a}}_{F} \overline{\boldsymbol{a}}_{F}^{T}}{f}-\frac{\overline{\boldsymbol{a}}^{T}}{m}\right| \leq \sqrt{\frac{A \log \left(\frac{1}{\epsilon}\right)}{f}}\right) & \geq 1-\epsilon,  \tag{A2}\\
\mathbb{P}\left(\left\|\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}-\frac{\overline{\boldsymbol{a}} \boldsymbol{V}}{m}\right\| \leq \sqrt{\frac{B k \log \left(\frac{k}{\epsilon}\right)}{f}}\right) & \geq 1-\epsilon, \tag{A3}
\end{align*}
$$

where $A, B$ are constants independent of $k, f, m, \epsilon$.
Now we would show that $\frac{\boldsymbol{R}^{T} Q^{T} Q \boldsymbol{Q}}{m}$ is close to $\frac{\boldsymbol{R}^{T} Q_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}$ :

## Lemma 1

$$
\begin{equation*}
\mathbb{P}\left(\left\|\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}-\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right\| \leq \sqrt{\frac{C k \log \left(\frac{k}{\epsilon}\right)}{f}}\right) \geq 1-\epsilon . \tag{A4}
\end{equation*}
$$

To prove this, we would first introduce a matrix analog of the well-known Chernoff bound, the proof of which can be found in Tropp (2012):

Lemma 2 Let $\mathcal{X} \in \mathbb{R}^{k \times k}$ be a finite set of positive-semidefinite matrices, and suppose that

$$
\max _{\boldsymbol{X} \in \mathcal{X}} \lambda_{\max }(\boldsymbol{X}) \leq D,
$$

where $\lambda_{\min } / \lambda_{\max }$ is the minimum/maximum eigenvalue function. Sample $\left\{\boldsymbol{X}_{1}, \ldots, \boldsymbol{X}_{\ell}\right\}$ uniformly at random without replacement. Compute:

$$
\mu_{\min }:=\ell \cdot \lambda_{\min }\left(\mathbb{E} \boldsymbol{X}_{1}\right) \quad \mu_{\max }:=\ell \cdot \lambda_{\max }\left(\mathbb{E} \boldsymbol{X}_{1}\right) .
$$

Then,

$$
\begin{aligned}
& \mathbb{P}\left\{\lambda_{\min }\left(\sum_{j} \boldsymbol{X}_{j}\right) \leq(1-\delta) \mu_{\min }\right\} \leq k \cdot \exp \left(\frac{-\delta^{2} \mu_{\min }}{4 D}\right), \quad \text { for } \delta \in[0,1), \\
& \mathbb{P}\left\{\lambda_{\max }\left(\sum_{j} \boldsymbol{X}_{j}\right) \leq(1+\delta) \mu_{\max }\right\} \leq k \cdot \exp \left(\frac{-\delta^{2} \mu_{\max }}{4 D}\right), \quad \text { for } \delta \geq 0 .
\end{aligned}
$$

Using this lemma, we would proceed with the proof of Lemma 1 .
Proof of Lemma 1: First, we note that

$$
\begin{aligned}
\boldsymbol{Q}^{T} \boldsymbol{Q} & =\sum_{i=1}^{m} \boldsymbol{q}_{i}^{T} \boldsymbol{q}_{i} \\
\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} & =\sum_{i \in F} \boldsymbol{q}_{i}^{T} \boldsymbol{q}_{i}
\end{aligned}
$$

where $\boldsymbol{q}_{i}^{T} \boldsymbol{q}_{i} \in \mathbb{R}^{k \times k}$ rank-one positive semi-definite matrices. Therefore, we can take $\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}$ as a random sample of size $f$ from the set $\mathcal{X}=\left\{\boldsymbol{q}_{i}^{T} \boldsymbol{q}_{i}\right\}_{i=1, \ldots, m}$, which satisfies the conditions in Lemma 22 with $D=O(k)$. Furthermore, with $\mathcal{X}$, we observe that we have $\mathbb{E} \boldsymbol{X}_{1}=\frac{Q^{T} \boldsymbol{Q}}{m}=\boldsymbol{I}_{k}$, so we have

$$
\lambda_{\min }\left(\mathbb{E} \boldsymbol{X}_{1}\right)=\lambda_{\max }\left(\mathbb{E} \boldsymbol{X}_{1}\right)=1
$$

Therefore, we apply Lemma 2 to $\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}$ and obtain

$$
\begin{aligned}
& \mathbb{P}\left\{\lambda_{\min }\left(\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}\right) \leq(1-\delta) f\right\} \leq k \cdot \exp \left(\frac{-\delta^{2} f}{k D^{\prime}}\right), \\
& \mathbb{P}\left\{\lambda_{\max }\left(\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}\right) \geq(1+\delta) f\right\} \leq k \cdot \exp \left(\frac{-\delta^{2} f}{k D^{\prime}}\right),
\end{aligned}
$$

where we set $D=\frac{k D^{\prime}}{4}$ with $D^{\prime}=O(1)$. Some rearrangement gives:

$$
\begin{equation*}
\mathbb{P}\left\{\lambda_{\min }\left(\frac{\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}}{f}\right) \geq 1-\sqrt{\frac{k D^{\prime} \log \left(\frac{2 k}{\epsilon}\right)}{f}} \text { and } \lambda_{\max }\left(\frac{\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}}{f}\right) \leq 1+\sqrt{\frac{k D^{\prime} \log \left(\frac{2 k}{\epsilon}\right)}{f}}\right\} \geq 1-\epsilon \tag{A5}
\end{equation*}
$$

Now since $\frac{\boldsymbol{Q}^{T} \boldsymbol{Q}}{m}=\boldsymbol{I}_{k}$, we have

$$
\begin{equation*}
\lambda_{\min }\left(\frac{\boldsymbol{Q}^{T} \boldsymbol{Q}}{m}\right)=\lambda_{\max }\left(\frac{\boldsymbol{Q}^{T} \boldsymbol{Q}}{m}\right)=1 \tag{A6}
\end{equation*}
$$

Combining equation (A6) and (A5) gives

$$
\begin{equation*}
\mathbb{P}\left\{\left\|\frac{\boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F}}{f}-\frac{\boldsymbol{Q}^{T} \boldsymbol{Q}}{m}\right\| \leq \sqrt{\frac{k D^{\prime} \log \left(\frac{2 k}{\epsilon}\right)}{f}}\right\} \geq 1-\epsilon \tag{A7}
\end{equation*}
$$

Then, we have

$$
\begin{equation*}
\mathbb{P}\left\{\left\|\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}-\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}\right\| \leq\|\boldsymbol{R}\|^{2} \sqrt{\frac{k D^{\prime} \log \left(\frac{2 k}{\epsilon}\right)}{f}}\right\} \geq 1-\epsilon . \tag{A8}
\end{equation*}
$$

Taking $C=D^{\prime}\|\boldsymbol{R}\|^{4} \log (2)$ gives the required result. (Note $\|\boldsymbol{R}\|=O(1)$ as we setup the QR decomposition to have $\boldsymbol{Q}^{T} \boldsymbol{Q}=O(m)$ ).

With equations (A2), (A3), and (A4), we are now ready to bound $\alpha(\boldsymbol{s})$ and $\alpha^{F}(\boldsymbol{s})$. We first introduce another lemma from matrix perturbation theory (for proof, see e.g. Stewart (1990)).

Lemma 3 Let $\boldsymbol{A}, \boldsymbol{B}$ be invertible matrices and let $\boldsymbol{B}=\boldsymbol{A}+\boldsymbol{\Delta}$. Then, we have

$$
\begin{equation*}
\left\|\boldsymbol{A}^{-1}-\boldsymbol{B}^{-1}\right\| \leq\left\|\boldsymbol{A}^{-1}\right\|\left\|\boldsymbol{B}^{-1}\right\|\|\boldsymbol{\Delta}\| . \tag{A9}
\end{equation*}
$$

Then, we have

$$
\begin{aligned}
\left\|\alpha(s)-\alpha^{F}(s)\right\| & \leq\left|\frac{\overline{\boldsymbol{a}}^{T}}{m}-\frac{\overline{\boldsymbol{a}}_{F} \overline{\boldsymbol{a}}_{F}^{T}}{f}\right| \\
& +\left|\frac{\overline{\boldsymbol{a}} \boldsymbol{V}}{m}\left(\frac{\boldsymbol{I}_{k}}{m \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}\right)^{-1} \frac{\boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}}{m}-\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}\left(\frac{\boldsymbol{I}_{k}}{f \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right)^{-1} \frac{\boldsymbol{V}_{F}^{T} \overline{\boldsymbol{a}}_{F}^{T}}{f}\right| .
\end{aligned}
$$

Using (A2) and triangle inequality, we have

$$
\begin{aligned}
& \leq \sqrt{\frac{A \log \left(\frac{1}{\epsilon}\right)}{f}}+\left|\left(\frac{\overline{\boldsymbol{a}} \boldsymbol{V}}{m}-\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}\right)\left(\frac{\boldsymbol{I}_{k}}{m \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}\right)^{-1} \frac{\boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}}{m}\right| \\
& +\left|\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}\left(\left(\frac{\boldsymbol{I}_{k}}{m \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}\right)^{-1}-\left(\frac{\boldsymbol{I}_{k}}{f \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right)^{-1}\right) \frac{\boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}}{m}\right| \\
& +\left|\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}\left(\frac{\boldsymbol{I}_{k}}{f \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right)^{-1}\left(\frac{\boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}}{m}-\frac{\boldsymbol{V}_{F}^{T} \overline{\boldsymbol{a}}_{F}^{T}}{f}\right)\right| .
\end{aligned}
$$

Using (A3) and Lemma 3, we have

$$
\begin{aligned}
& \leq \sqrt{\frac{A \log \left(\frac{1}{\epsilon}\right)}{f}}+\sqrt{\frac{B^{\prime} k \log \left(\frac{k}{\epsilon}\right)}{f}} \\
& +D^{\prime}\left|\frac{\overline{\boldsymbol{a}}_{F} \boldsymbol{V}_{F}}{f}\left(\frac{\boldsymbol{I}_{k}}{m \gamma}-\frac{\boldsymbol{I}_{k}}{f \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}-\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right) \frac{\boldsymbol{V}^{T} \overline{\boldsymbol{a}}^{T}}{m}\right| \\
& +\sqrt{\frac{B^{\prime \prime} k \log \left(\frac{k}{\epsilon}\right)}{f}}
\end{aligned}
$$

for some $O(1)$ constant $D^{\prime}$. Note that $\left\|\frac{I_{k}}{m \gamma}-\frac{I_{k}}{f \gamma}\right\| \leq \frac{M}{f}$, so using (A4), we obtain

$$
\begin{equation*}
\mathbb{P}\left(\left\|\frac{\boldsymbol{I}_{k}}{m \gamma}-\frac{\boldsymbol{I}_{k}}{f \gamma}+\frac{\boldsymbol{R}^{T} \boldsymbol{Q}^{T} \boldsymbol{Q} \boldsymbol{R}}{m}-\frac{\boldsymbol{R}^{T} \boldsymbol{Q}_{F}^{T} \boldsymbol{Q}_{F} \boldsymbol{R}}{f}\right\| \leq \sqrt{\frac{C^{\prime} k \log \left(\frac{k}{\epsilon}\right)}{f}}\right) \geq 1-\epsilon, \tag{A10}
\end{equation*}
$$

for some new $O(1)$ constant $C^{\prime}$. Then, using (A10), we have

$$
\leq \sqrt{\frac{A \log \left(\frac{1}{\epsilon}\right)}{f}}+\sqrt{\frac{B^{\prime} k \log \left(\frac{k}{\epsilon}\right)}{f}}+\sqrt{\frac{C^{\prime \prime} k \log \left(\frac{k}{\epsilon}\right)}{f}}+\sqrt{\frac{B^{\prime \prime} k \log \left(\frac{k}{\epsilon}\right)}{f}} .
$$

Since $k \geq 1$, we have

$$
\leq \sqrt{\frac{M k \log \left(\frac{k}{\epsilon}\right)}{f}}
$$

For some $O(1)$ constant $M$ with probability at least $1-4 \epsilon$. Therefore, taking $\epsilon^{\prime}=\frac{\epsilon}{4}$ and absorbing the extra constant into $M$ gives the requires result.

## B. Proof of Theorem 3

We first prove a Hoeffding-type bound as follows

Proposition 3 Let $X_{1}, \ldots, X_{n}$ be independent (but not necessarily identically distributed) random variables which satisfy

$$
\mathbb{P}\left(\left|X_{i}-a_{i}\right| \geq t\right) \leq \exp \left(\frac{-t^{2}}{\sigma_{i}^{2}}\right)
$$

Then, we have

$$
\begin{equation*}
\mathbb{P}\left(\left|\frac{\sum_{i=1}^{n} X_{i}-a_{i}}{n}\right| \geq t\right) \leq 4 \exp \left(-\frac{n^{2} t^{2}}{4 \sum_{i=1}^{n} \sigma_{i}^{2}}\right) . \tag{A11}
\end{equation*}
$$

We note that such statement differs from Hoeffding's inequality as we do not require $\mathbb{E}\left[X_{i}\right]=a_{i}$. The proof is as follows.

Proof of Proposition 3: We first introduce a lemma known as Chernoff's bounding method (proven in Chernoff (1952)):

Lemma 4 Let $Z$ be a random variable on $\mathbb{R}$. Then for all $t>0$, we have

$$
\mathbb{P}(Z \geq t) \leq \inf _{s>0}\left[e^{-s t} M_{Z}(s)\right],
$$

where $M_{Z}(s)$ is the moment generating function of $Z$.
Let $Y_{i}=X_{i}-a_{i}$. Then, we have

$$
\begin{equation*}
\mathbb{P}\left(\sum_{i=1}^{n} Y_{i} \geq t\right) \leq \min _{s>0} e^{-s t} \prod_{i=1}^{n} \mathbb{E}\left[e^{s Y_{i}}\right] . \tag{A12}
\end{equation*}
$$

We aim to bound $\mathbb{E}\left[e^{s Y_{i}}\right]$, subject to $\mathbb{P}\left(\left|Y_{i}\right| \geq t\right) \leq \exp \left(\frac{-t^{2}}{\sigma_{i}^{2}}\right)$. Since $e^{s t}$ is an increasing function of $t, \mathbb{E}\left[e^{s Y_{i}}\right]$ is maximized when $\mathbb{P}\left(Y_{i} \geq t\right)=\exp \left(\frac{-t^{2}}{\sigma_{i}^{2}}\right)$, which results in a probability distribution function of $f_{Y_{i}}(y)=2 \frac{t}{\sigma_{i}^{2}} \exp \left(\frac{-t^{2}}{\sigma_{i}^{2}}\right)$ for $Y_{i}$. Then, we have

$$
\begin{aligned}
\mathbb{E}\left[e^{s Y_{i}}\right] & \leq \int_{0}^{\infty} 2 \frac{t}{\sigma_{i}^{2}} \exp \left(s t-\frac{t^{2}}{\sigma_{i}^{2}}\right) \mathrm{d} t \\
& =\int_{0}^{\infty} 2 \frac{t}{\sigma_{i}^{2}} \exp \left(-\left(\frac{t}{\sigma_{i}}-\frac{s \sigma_{i}}{2}\right)^{2}+\frac{s^{2} \sigma_{i}^{2}}{4}\right) \mathrm{d} t \\
& =2 \exp \left(\frac{s^{2} \sigma_{i}^{2}}{4}\right) \int_{0}^{\infty} t \exp \left(-\left(t-\frac{s \sigma_{i}}{2}\right)^{2}\right) \mathrm{d} t \\
& =\exp \left(\frac{s^{2} \sigma_{i}^{2}}{4}\right)\left(\exp \left(-\frac{s^{2} \sigma_{i}^{2}}{4}\right)+\sqrt{\pi} \frac{s \sigma_{i}}{2}\left(\operatorname{erf}\left(\frac{s \sigma_{i}}{2}\right)+1\right)\right)
\end{aligned}
$$

where erf is the standard error function. As the error function is upper bounded by 1 , the last expression is less than or equal to:

$$
\begin{aligned}
& \leq 1+\sqrt{\pi} s \sigma_{i} \exp \left(\frac{s^{2} \sigma_{i}^{2}}{4}\right) \\
& \leq 2 \exp \left(s^{2} \sigma_{i}^{2}\right)
\end{aligned}
$$

We then substitute this result into (A12) and obtain

$$
\mathbb{P}\left(\sum_{i=1}^{n} Y_{i} \geq t\right) \leq \min _{s>0} 2 \exp \left(-s t+s^{2} \sum_{i=1}^{n} \sigma_{i}^{2}\right)
$$

Note that this is minimized at $s=\frac{t}{2 \sqrt{\sum_{i=1}^{n} \sigma_{i}^{2}}}$, so we have

$$
\leq 2 \exp \left(-\frac{t^{2}}{4 \sum_{i=1}^{n} \sigma_{i}^{2}}\right)
$$

Therefore,

$$
\mathbb{P}\left(\frac{\sum_{i=1}^{n} Y_{i}}{n} \geq t\right) \leq 2 \exp \left(-\frac{n^{2} t^{2}}{4 \sum_{i=1}^{n} \sigma_{i}^{2}}\right) .
$$

By applying the previous derivation to $-Y_{1}, \ldots,-Y_{n}$, we obtain

$$
\mathbb{P}\left(\frac{\sum_{i=1}^{n} Y_{i}}{n} \leq-t\right) \leq 2 \exp \left(-\frac{n^{2} t^{2}}{4 \sum_{i=1}^{n} \sigma_{i}^{2}}\right) .
$$

Combining the two results completes the proof.
Using Proposition 3, we bound the difference between $\tilde{c}_{G}^{F}\left(s_{t}\right)$ and $c\left(s_{t}\right)$. We have

$$
\begin{aligned}
\left|\tilde{c}_{G}^{F}\left(s_{t}\right)-c\left(s_{t}\right)\right| & =\left|\frac{1}{n} \sum_{i=1}^{n} \alpha_{i}(s)-\frac{1}{g} \sum_{i \in G} \alpha_{i}^{F}(s)\right| \\
& \leq\left|\frac{1}{n} \sum_{i=1}^{n} \alpha_{i}(\boldsymbol{s})-\frac{1}{g} \sum_{i \in G} \alpha_{i}(\boldsymbol{s})\right|+\left|\frac{1}{g} \sum_{i \in G} \alpha_{i}(\boldsymbol{s})-\frac{1}{g} \sum_{i \in G} \alpha_{i}^{F}(\boldsymbol{s})\right|
\end{aligned}
$$

The first term can be seen as the tail bound for a random sample of size $g$ chosen without replacement from the finite set $\left\{\alpha_{i}(s)\right\}_{i=1, \ldots, n}$. Thus, we can apply Hoeffding's theorem in Proposition 2, and obtain that with probability at least $1-\epsilon$

$$
\begin{equation*}
\left|\frac{1}{n} \sum_{i=1}^{n} \alpha_{i}(s)-\frac{1}{g} \sum_{i \in G} \alpha_{i}(s)\right| \leq \sqrt{\frac{M \log \left(\frac{1}{\epsilon}\right)}{g}} \tag{A13}
\end{equation*}
$$

Substituting (A13) into the expression above shows that with probability at least $1-\epsilon$ we have:

$$
\leq \sqrt{\frac{M \log \left(\frac{1}{\epsilon}\right)}{g}}+\left|\frac{1}{g} \sum_{i \in G} \alpha_{i}(s)-\frac{1}{g} \sum_{i \in G} \alpha_{i}^{F}(s)\right|
$$

Note that, for any fixed set $G$, for $i, j \in G$ with $i \neq j$, we have that $\alpha_{i}^{F}(\boldsymbol{s})$ and $\alpha_{j}^{F}(\boldsymbol{s})$ are independent (as we construct $F$ separately for every $i$ ). Furthermore, Theorem 2 can be inverted to read

$$
\mathbb{P}\left(\left|\alpha_{i}(\boldsymbol{s})-\alpha_{i}^{F}(\boldsymbol{s})\right| \geq t\right) \leq \exp \left(-\frac{f t^{2}}{M k \log (k)}\right)
$$

Therefore, $\alpha_{i}^{F}(\boldsymbol{s})$ satisfies the conditions of Proposition 3 with $X_{i}=\alpha_{i}^{F}(\boldsymbol{s})$ and parameters $\sigma_{i}^{2}=$ $\frac{M k \log (k)}{f}, a_{i}=\alpha_{i}(\boldsymbol{s})$. Then, applying Proposition 3 to the second term, we have

$$
\leq \sqrt{\frac{M \log \left(\frac{1}{\epsilon}\right)}{g}}+\sqrt{\frac{M^{\prime} k \log \left(\frac{k}{\epsilon}\right)}{f g}}
$$

with probability $1-2 \epsilon$. As $k, f \geq 1$, a loose bound is therefore

$$
\leq \sqrt{\frac{M^{\prime \prime} k \log \left(\frac{k}{\epsilon}\right)}{g}}
$$

with probability $1-2 \epsilon$. Therefore, taking $\epsilon^{\prime}=2 \epsilon$ and $A=M^{\prime \prime} \log (2)$ we have

$$
\begin{equation*}
\mathbb{P}\left(\left|\tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)-c\left(\boldsymbol{s}_{t}\right)\right| \leq \sqrt{\frac{A k \log \left(\frac{k}{\epsilon^{\prime}}\right)}{g}}\right) \geq 1-\epsilon^{\prime} \tag{A14}
\end{equation*}
$$

Through a similar derivation, we have

$$
\begin{equation*}
\mathbb{P}\left(\left\|\nabla \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)-\nabla c\left(\boldsymbol{s}_{t}\right)\right\|_{2} \leq \sqrt{\frac{B(p+k) \log \left(\frac{k}{\epsilon^{\prime}}\right)}{g}}\right) \geq 1-\epsilon^{\prime} \tag{A15}
\end{equation*}
$$

## C. Proof of Proposition 1

We would prove the contrapositive. Assume that $a=0$. We would calculate the second derivative of the cost function expression in 7. First, define $\boldsymbol{T}_{i j}=\frac{I}{\gamma}+\boldsymbol{W}_{i} \boldsymbol{K}_{j} \boldsymbol{W}_{i}$. Then our cost function can be rewritten as:

$$
c(\boldsymbol{s})=\frac{1}{n m} \sum_{i=1}^{n} \overline{\boldsymbol{a}}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}
$$

Then the second derivative can be easily calculated as:

$$
\frac{\partial c(\boldsymbol{s})}{\partial s_{k} \partial s_{l}}=2 \sum_{i=1}^{n} \overline{\boldsymbol{a}}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \boldsymbol{T}_{i k}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \boldsymbol{T}_{i l}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}
$$

Thus, for any vector $\boldsymbol{t} \in \mathbb{R}^{p}$, we have:

$$
\boldsymbol{t}^{T} \frac{\partial c(\boldsymbol{s})}{\partial s_{k} \partial s_{l}} \boldsymbol{t}=2 \sum_{i=1}^{n} \overline{\boldsymbol{a}}_{i}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1}\left(\sum_{k=1}^{p} t_{k} \boldsymbol{T}_{i k}\right)\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1}\left(\sum_{l=1}^{p} t_{l} \boldsymbol{T}_{i l}\right)\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}
$$

Now define $\boldsymbol{v}_{i}=\left(\sum_{l=1}^{p} t_{l} \boldsymbol{T}_{i l}\right)\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}$, then our expression becomes:

$$
\boldsymbol{t}^{T} \frac{\partial c(\boldsymbol{s})}{\partial s_{k} \partial s_{l}} \boldsymbol{t}=2 \sum_{i=1}^{n} \boldsymbol{v}_{i}^{T}\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \boldsymbol{v}_{i}
$$

Now since $a=0$, the Hessian cannot be positive definite for all $s$ (as if it was, then $a>0$ ). Then there exist $s^{0}, s^{1} \in S_{k}^{p}$ such that for $t=s^{0}-s^{1}$, we have:

$$
\left.\boldsymbol{t}^{T} \frac{\partial c(\boldsymbol{s})}{\partial s_{k} \partial s_{l}}\right|_{\boldsymbol{s}=\boldsymbol{s}^{0}} \boldsymbol{t}=2 \sum_{i=1}^{n} \boldsymbol{v}_{i}^{T}\left(\sum_{j=1}^{p} s_{j}^{0} \boldsymbol{T}_{i j}\right)^{-1} \boldsymbol{v}_{i}=0
$$

and $\boldsymbol{t}$ Note that for any $\gamma>0, \boldsymbol{T}_{i j}$ is positive definite, so $\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}$ is positive definite for any $s \in S_{k}^{p}$ and all $i$. Therefore if $a=0$, we must have $\boldsymbol{v}_{i}=0$ for all $i$, which means that:

$$
\left(\sum_{l=1}^{p} t_{l} \boldsymbol{T}_{i l}\right)\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}=\mathbf{0}
$$

for all $i$. Now note that we have:

$$
c\left(\boldsymbol{s}^{1}\right)=c\left(\boldsymbol{s}^{0}\right)+\nabla c\left(\boldsymbol{s}^{0}\right)^{T} \boldsymbol{t}+\left.\boldsymbol{t}^{T} \frac{\partial c(\boldsymbol{s})}{\partial s_{k} \partial s_{l}}\right|_{\boldsymbol{s}=\boldsymbol{s}^{0}} \boldsymbol{t}+\cdots
$$

We have that:

$$
\nabla c\left(\boldsymbol{s}^{0}\right)^{T} \boldsymbol{t}=\sum_{i=1}^{n}\left(\sum_{l=1}^{p} t_{l} \boldsymbol{T}_{i l}\right)\left(\sum_{j=1}^{p} s_{j} \boldsymbol{T}_{i j}\right)^{-1} \overline{\boldsymbol{a}}_{i}^{T}=0
$$

And similarly, we have $\left.\boldsymbol{t}^{T} \frac{\partial c(s)}{\partial s_{k} \partial s_{l}}\right|_{s=s^{0}} \boldsymbol{t}=0$ and all higher derivatives being 0 .
Thus, we have:

$$
c\left(\boldsymbol{s}^{1}\right)=c\left(\boldsymbol{s}^{0}\right)
$$

Therefore, if we have $a=0$, then we must have equality on the cost function for two feasible solutions. Thus, if the cost function is not degenerate for feasible solutions, then $a>0$.

## D. Proof of Theorem 4

In this section, we provide the proof of Theorem 4. We first note that OptComplete is a specific implementation of the outer approximation algorithm, which Fletcher and Leyffer (1994) have shown to always terminate in finite number of steps $C$. Thus, given that we assumed the problem is feasible, for OptComplete to not return an optimal solution, it would have to cut it off during the course of its execution. Let $s^{*}$ be an optimal solution for Problem (4). Let $s_{t}$ be an optimal solution at the $t$-th iteration of OptComplete, $t \in[C]$. The cutting plane constraint for OptComplete at the point of an optimal solution $\boldsymbol{s}^{*}$ is

$$
\eta_{t} \geq \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)+\nabla \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)^{T}\left(s^{*}-\boldsymbol{s}_{t}\right)
$$

If $c\left(s^{*}\right)<\eta_{t}$, then $s^{*}$ will be cut off, and OptComplete will not find $s^{*}$. Applying the definition of the convexity parameter (16) and letting $\left\|s^{*}-s_{t}\right\|=\theta_{t}$ (noting that $\theta_{t} \geq 1$ ) we obtain

$$
\begin{equation*}
c\left(\boldsymbol{s}^{*}\right) \geq c\left(\boldsymbol{s}_{t}\right)+\nabla c\left(\boldsymbol{s}_{t}\right)^{T}\left(\boldsymbol{s}^{*}-\boldsymbol{s}_{t}\right)+\frac{\theta_{t}^{2} a^{2}}{2} \tag{A16}
\end{equation*}
$$

Therefore, if

$$
c\left(\boldsymbol{s}_{t}\right)+\nabla c\left(\boldsymbol{s}_{t}\right)^{T}\left(\boldsymbol{s}^{*}-\boldsymbol{s}_{t}\right)+\frac{\theta_{t}^{2} a^{2}}{2} \leq c\left(\boldsymbol{s}^{*}\right)<\tilde{c}_{r}\left(\boldsymbol{s}_{t}\right)+\nabla \tilde{c}_{r}\left(\boldsymbol{s}_{t}\right)^{T}\left(\boldsymbol{s}^{*}-\boldsymbol{s}_{t}\right)
$$

or equivalently if

$$
\begin{equation*}
\zeta_{t}:=\left[\tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)-c\left(\boldsymbol{s}_{t}\right)\right]+\left[\nabla \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)-\nabla c\left(\boldsymbol{s}_{t}\right)\right]^{T}\left(\boldsymbol{s}^{*}-\boldsymbol{s}_{t}\right)>\frac{\theta_{t}^{2} a^{2}}{2} \tag{A17}
\end{equation*}
$$

then OptComplete will not find $s^{*}$. Therefore, for OptComplete to succeed, $\zeta_{t}$ should satisfy $\zeta_{t} \leq$ $\theta_{t}^{2} a^{2} / 2$ for all $t \in[C]$.

Let $S=$ the event that OptComplete succeeds in finding $\boldsymbol{s}^{*}$. Then,

$$
\mathbb{P}(S) \geq \mathbb{P}\left(\zeta_{t} \leq \frac{\theta_{t}^{2} a^{2}}{2} t \in[C]\right)
$$

Since at each step of OptComplete we randomly sample $r$ new rows and $s$ new columns, the events $\zeta_{t} \leq \frac{\theta_{t}^{2} a^{2}}{2}$ are independent for different $t \in[C]$, and hence

$$
\begin{equation*}
\mathbb{P}(S) \geq \prod_{t=1}^{C}\left(1-\mathbb{P}\left(\zeta_{t}>\frac{\theta_{t}^{2} a^{2}}{2}\right)\right) \tag{A18}
\end{equation*}
$$

Therefore, we focus on calculating $\mathbb{P}\left(\zeta_{t} \geq \theta_{t}^{2} a^{2} / 2\right)$. Since

$$
\zeta_{t}:=\left[\tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)-c\left(\boldsymbol{s}_{t}\right)\right]+\left[\nabla \tilde{c}_{G}^{F}\left(\boldsymbol{s}_{t}\right)-\nabla c\left(\boldsymbol{s}_{t}\right)\right]^{T}\left(\boldsymbol{s}^{*}-\boldsymbol{s}_{t}\right),
$$

using Theorem 3, we can thus provide the following bound for deviation of $\zeta_{t}$, for some constant $C$ :

$$
\begin{equation*}
\mathbb{P}\left(\zeta_{t} \leq \theta_{t} \sqrt{\frac{C(p+k) \log \left(\frac{k}{\epsilon}\right)}{g}}\right) \geq 1-\epsilon, \tag{A19}
\end{equation*}
$$

where $\theta_{t}=\left\|s^{*}-s_{t}\right\|$. Then we can invert this to calculate $\mathbb{P}\left(\zeta_{t} \geq \theta_{t}^{2} a^{2} / 2\right)$

$$
\begin{equation*}
\mathbb{P}\left(\zeta_{t} \geq \theta_{t}^{2} a^{2} / 2\right) \leq k \exp \left(-\frac{a^{4} g}{4 C(p+k) \theta_{t}^{2}}\right) \leq k \exp \left(-\frac{D a^{4} g}{(p+k)}\right) \tag{A20}
\end{equation*}
$$

taking $D=\frac{1}{4 C}$, and noting that $\theta_{t} \geq 1$. Then, we substitute (A20) into (A18) to obtain

$$
\begin{aligned}
\mathbb{P}(S) & \geq\left(1-k \exp \left(-\frac{D a^{4} g}{(p+k)}\right)\right)^{C} \\
& \geq 1-k C \exp \left(-\frac{D a^{4} g}{(p+k)}\right)
\end{aligned}
$$

completing the proof.

## E. List of Features Used in the Netflix Problem

- 24 Indicator Variables for Genres: Action, Adventure, Animation, Biography, Comedy, Crime, Documentary, Drama, Family, Fantasy, Film Noir, History, Horror, Music, Musical, Mystery, Romance, Sci-Fi, Short, Sport, Superhero, Thriller, War, Western
- 5 Indicator Variables for Release Date: Within last 10 years, Between 10-20 years, Between 20-30 years, Between 30-40 years, Between 40-50 Years
- 6 Indicator Variables for Top Actors/Actresses defined by their Influence Score at time of release: Top 100 Actors, Top 100 Actresses, Top 250 Actors, Top 250 Actresses, Top 1000 Actors, Top 1000 Actresses
- IMDB Rating
- Number of Reviews
- Total Production Budget
- Total Runtime
- Total Box Office Revenue
- Indicator Variable for whether it is US produced
- 11 Indicator Variables for Month of Year Released (January removed to prevent multicollinear-
ity)
- Number of Original Music Score
- Number of Male Actors
- Number of Female Factors
- 3 Indicator Variables for Film Language: English, French, Japanese
- Constant


## References

Beck A, Teboulle M (2009) A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences 2(1):183-202.

Bennett J, Lanning S, et al. (2007) The netflix prize. KDD Cup and Workshop 2007 (Citeseer).

Bertsimas D, Copenhaver MS (2018) Characterization of the equivalence of robustification and regularization in linear, median, and matrix regression. European Journal of Operations Research 270:931-942.

Bertsimas D, van Parys B (2020) Sparse high-dimensional regression: Exact scalable algorithms and phase transitions. Annals of Statistics 48(1):300-323.

Boucheron S, Lugosi G, Massart P (2013) Concentration inequalities: A nonasymptotic theory of independence (Oxford university press).

Boyd S, El Ghaoui L, Feron E, Balakrishnan V (1994) Linear matrix inequalities in system and control theory, volume 15 (SIAM).

Candes EJ, Plan Y (2010) Matrix completion with noise. Proceedings of the IEEE 98(6):925-936.

Candès EJ, Tao T (2010) The power of convex relaxation: Near-optimal matrix completion. IEEE Transactions on Information Theory 56(5):2053-2080.

Chen Y, Bhojanapalli S, Sanghavi S, Ward R (2014a) Coherent matrix completion. International Conference on Machine Learning, 674-682.

Chen Y, Jalali A, Sanghavi S, Xu H (2014b) Clustering partially observed graphs via convex optimization. The Journal of Machine Learning Research 15(1):2213-2238.

Chernoff H (1952) A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. The Annals of Mathematical Statistics 23(4):493-507.

Chiang KY, Hsieh CJ, Dhillon IS (2015) Matrix completion with noisy side information. Advances in Neural Information Processing Systems, 3447-3455.

Chiang KY, Hsieh CJ, Natarajan N, Dhillon IS, Tewari A (2014) Prediction and clustering in signed networks: a local to global perspective. The Journal of Machine Learning Research 15(1):1177-1213.

Dhillon P, Lu Y, Foster DP, Ungar L (2013) New subsampling algorithms for fast least squares regression. Advances in Neural Information Processing Systems, 360-368.

Duran MA, Grossmann IE (1986) An outer-approximation algorithm for a class of mixed-integer nonlinear programs. Mathematical Programming 36(3):307-339.

Farhat MR, Shapiro BJ, Kieser KJ, Sultana R, Jacobson KR, Victor TC, Warren RM, Streicher EM, Calver A, Sloutsky A, et al. (2013) Genomic analysis identifies targets of convergent positive selection in drug-resistant mycobacterium tuberculosis. Nature Genetics 45(10):1183.

Fletcher R, Leyffer S (1994) Solving mixed integer nonlinear programs by outer approximation. Mathematical programming 66(1-3):327-349.

Hastie T, Mazumder R, Lee JD, Zadeh R (2015) Matrix completion and low-rank svd via fast alternating least squares. The Journal of Machine Learning Research 16(1):3367-3402.

Jain P, Dhillon IS (2013) Provable inductive matrix completion. arXiv preprint arXiv:1306.0626 .
Jain P, Meka R, Dhillon IS (2010) Guaranteed rank minimization via singular value projection. Advances in Neural Information Processing Systems, 937-945.

Ji H, Liu C, Shen Z, Xu Y (2010) Robust video denoising using low rank matrix completion. Computer Society Conference on Computer Vision and Pattern Recognition (IEEE).

Keshavan RH, Oh S, Montanari A (2009) Matrix completion from a few entries. Information Theory, 2009. ISIT 2009. IEEE International Symposium on, 324-328 (IEEE).

Koren Y, Bell R, Volinsky C (2009) Matrix factorization techniques for recommender systems. Computer 30-37.

Lu J, Liang G, Sun J, Bi J (2016) A sparse interactive model for matrix completion with side information. Advances in Neural Information Processing Ssystems, 4071-4079.

Lubin M, Yamangil E, Bent R, Vielma JP (2016) Extended formulations in mixed-integer convex programming. International Conference on Integer Programming and Combinatorial Optimization, 102-113 (Springer).

Mazumder R, Hastie T, Tibshirani R (2010) Spectral regularization algorithms for learning large incomplete matrices. Journal of Machine Learning Research 11(Aug):2287-2322.

Natarajan N, Dhillon IS (2014) Inductive matrix completion for predicting gene-disease associations. Bioinformatics 30(12):160-168.

Nazarov I, Shirokikh B, Burkina M, Fedonin G, Panov M (2018) Sparse group inductive matrix completion. arXiv preprint arXiv:1804.10653 .

Negahban S, Wainwright MJ (2012) Restricted strong convexity and weighted matrix completion: Optimal bounds with noise. Journal of Machine Learning Research 13(May):1665-1697.

Recht B, Ré C (2013) Parallel stochastic gradient algorithms for large-scale matrix completion. Mathematical Programming Computation 5(2):201-226.

Shah V, Rao N, Ding W (2017) Matrix factorization with side and higher order information. Stat 1050:4.

Si S, Chiang KY, Hsieh CJ, Rao N, Dhillon IS (2016) Goal-directed inductive matrix completion. Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 11651174 (ACM).

Soni A, Chevalier T, Jain S (2016) Noisy inductive matrix completion under sparse factor models. arXiv preprint arXiv:1609.03958 .

Stewart GW (1990) Matrix perturbation theory .

Tanner J, Wei K (2013) Normalized iterative hard thresholding for matrix completion. SIAM Journal on Scientific Computing 35(5):S104-S125.

Tropp JA (2012) User-friendly tail bounds for sums of random matrices. Foundations of computational mathematics 12(4):389-434.

Woodbury MA (1949) The stability of out-input matrices. Chicago, Ill .

Xu M, Jin R, Zhou ZH (2013) Speedup matrix completion with side information: Application to multi-label learning. Advances in Neural Information Processing Systems, 2301-2309.

