
Bipartite Correlation Clustering – Maximizing Agreements

Megasthenis Asteris

The University of Texas at Austin

Dimitris Papailiopoulos

University of California, Berkeley

Anastasios Kyrillidis

The University of Texas at Austin

Alexandros G. Dimakis

The University of Texas at Austin

Abstract

In *Bipartite Correlation Clustering* (BCC) we are given a complete *bipartite* graph G with ‘+’ and ‘-’ edges, and we seek a vertex clustering that maximizes the number of *agreements*: the number of all ‘+’ edges within clusters plus all ‘-’ edges cut across clusters. BCC is known to be NP-hard [5].

We present a novel approximation algorithm for k -BCC, a variant of BCC with an upper bound k on the number of clusters. Our algorithm outputs a k -clustering that provably achieves a number of agreements within a multiplicative $(1 - \delta)$ -factor from the optimal, for any desired accuracy δ . It relies on solving a combinatorially constrained bilinear maximization on the bi-adjacency matrix of G . It runs in time exponential in k and $1/\delta$, but linear in the size of the input.

Further, we show that, in the (unconstrained) BCC setting, an $(1 - \delta)$ -approximation can be achieved by $O(\delta^{-1})$ clusters regardless of the size of the graph. In turn, our k -BCC algorithm implies an Efficient PTAS for the BCC objective of maximizing agreements.

1 Introduction

Correlation Clustering (CC) [7] considers the task of partitioning a set of objects into clusters based on their pairwise relationships. It arises naturally in several areas such as in network monitoring [1], document clustering [7] and textual similarity for data integration [10]. In its simplest form, objects are represented as vertices in a complete graph whose edges are labeled ‘+’ or ‘-’ to encode *similarity* or *dissimilarity* among vertices, respectively. The objective is to compute a vertex partition that maximizes the number of *agreements* with the underlying graph, *i.e.*, the total

number of ‘+’ edges in the interior of clusters plus the number of ‘-’ edges across clusters. The number of output clusters is itself an optimization variable –not part of the input. It may be meaningful, however, to restrict the number of output clusters to be at most k . The constrained version is known as k -CC and similarly to the unconstrained problem, it is NP-hard [7, 13, 15]. A significant volume of work has focused on approximately solving the problem of maximizing the number of agreements (MAXAGREE), or the equivalent –yet more challenging in terms of approximation– objective of minimizing the number of *disagreements* (MINDISAGREE) [7, 11, 9, 3, 19, 4].

Bipartite Correlation Clustering (BCC) is a natural variant of CC on bipartite graphs. Given a complete bipartite graph $G = (U, V, E)$ with edges labeled ‘+’ or ‘-’, the objective is once again to compute a clustering of the vertices that maximizes the number of agreements with the labeled pairs. BCC is a special case of the *incomplete* CC problem [9], where only a subset of vertices are connected and non-adjacent vertices do not affect the objective function. The output clusters may contain vertices from either one or both sides of the graph. Finally, we can define the k -BCC variant that enforces an upper bound on the number of output clusters, similar to k -CC for CC.

The task of clustering the vertices of a bipartite graph is common across many areas of machine learning. Applications include recommendation systems [18, 20], where analyzing the structure of a large sets of pairwise interactions (*e.g.*, among users and products) allows useful predictions about future interactions, gene expression data analysis [5, 16] and graph partitioning problems in data mining [12, 21].

Despite the practical interest in bipartite graphs, there is limited work on BCC and it is focused on the theoretically more challenging MINDISAGREE objective: [5] established an 11-approximation algorithm, while [2] achieved a 4-approximation, the currently best known guarantee. Algorithms for incomplete CC

[11, 9, 17] can be applied to BCC, but they do not leverage the structure of the bipartite graph. Moreover, existing approaches for incomplete CC rely on LP or SDP solvers which scale poorly.

Our contributions We develop a novel approximation algorithm for k -BCC with provable guarantees for the MAXAGREE objective. Further, we show that under an appropriate configuration, our algorithm yields an Efficient Polynomial Time Approximation Scheme (EPTAS¹) for the unconstrained BCC problem. Our contributions can be summarized as follows:

1. k -BCC: Given a bipartite graph $G = (U, V, E)$, a parameter k , and any constant accuracy parameter $\delta \in (0, 1)$, our algorithm computes a clustering of $U \cup V$ into at most k clusters and achieves a number of agreements that lies within a $(1 - \delta)$ -factor from the optimal. It runs in time exponential in k and δ^{-1} , but linear in the size of G .
2. BCC: In the unconstrained BCC setting, the optimal number of clusters may be anywhere from 1 to $|U| + |V|$. We show that if one is willing to settle for a $(1 - \delta)$ -approximation of the MAXAGREE objective, it suffices to use at most $O(\delta^{-1})$ clusters, regardless of the size of G . In turn, under an appropriate configuration, our k -BCC algorithm yields an EPTAS for the (unconstrained) BCC problem.
3. Our algorithm relies on formulating the k -BCC/MAXAGREE problem as a combinatorially constrained bilinear maximization

$$\max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \mathbf{B} \mathbf{Y}), \quad (1)$$

where \mathbf{B} is the bi-adjacency matrix of G , and \mathcal{X} , \mathcal{Y} are the sets of cluster assignment matrices for U and V , respectively. In Sec. 3, we briefly describe our approach for approximately solving (1) and its guarantees under more general constraints.

We note that our k -BCC algorithm and its guarantees can be extended to *incomplete*, but dense BCC instances, where the input G is *not* a complete bipartite graph, but $|E| = \Omega(|U| \cdot |V|)$. For simplicity, we restrict the description to the complete case. Finally, we supplement our theoretical findings with experimental results on synthetic and real datasets.

¹EPTAS refers to an algorithm that approximates the solution of an optimization problem within a multiplicative $(1 - \epsilon)$ -factor, for any constant $\epsilon \in (0, 1)$, and has complexity that scales arbitrarily in $1/\epsilon$, but as a constant order polynomial (independent of ϵ) in the input size n . EPTAS is more efficient than a PTAS; for example, a running time of $O(n^{1/\epsilon})$ is considered a PTAS, but not an EPTAS.

1.1 Related work

There is extensive literature on CC; see [8] for a list of references. Here, we focus on bipartite variants.

BCC was introduced by Amit in [5] along with an 11-approximation algorithm for the MINDISAGREE objective, based on a linear program (LP) with $O(|E|^3)$ constraints. In [2], Ailon et al. proposed two algorithms for the same objective: *i*) a deterministic 4-approximation based on an LP formulation and derandomization arguments by [19], and *ii*) a randomized 4-approximation combinatorial algorithm. The latter is computationally more efficient with complexity scaling linearly in the size of the input, compared to the LP that has $O((|V| + |U|)^3)$ constraints.

For the *incomplete* CC problem, which encompasses BCC as a special case, [11] provided an LP-based $O(\log n)$ -approximation for MINDISAGREE. A similar result was achieved by [9]. For the MAXAGREE objective, [2, 17] proposed an SDP relaxation, similar to that for MAX k -CUT, and achieved a 0.7666-approximation. We are not aware of any results explicitly on the MAXAGREE objective for either BCC or k -BCC. For comparison, in the k -CC setting [7] provided a simple 3-approximation algorithm for $k = 2$, while for $k \geq 2$ [13] provided a PTAS for MAXAGREE and one for MINDISAGREE. [15] improved on the latter utilizing approximations schemes to the Gale-Berlekamp switching game. Table 1 summarizes the aforementioned results.

Finally, we note that our algorithm relies on adapting ideas from [6] for approximately solving a combinatorially constrained quadratic maximization.

2 k -BCC as a Bilinear Maximization

We describe the k -BCC problem and show that it can be formulated as a combinatorially constrained bilinear maximization on the bi-adjacency matrix of the input graph. Our k -BCC algorithm relies on approximately solving that bilinear maximization.

Maximizing Agreements An instance of k -BCC consists of an undirected, complete, bipartite graph $G = (U, V, E)$ whose edges have binary ± 1 weights, and a parameter k . The objective, referred to as MAXAGREE[k], is to compute a clustering of the vertices into at most k clusters, such that the total number of positive edges in the interior of clusters plus the number of negative edges across clusters is maximized.

Let E^+ and E^- denote the sets of positive and negative edges, respectively. Also, let $m = |U|$ and $n = |V|$. The bipartite graph G can be represented by its weighted bi-adjacency matrix $\mathbf{B} \in \{\pm 1\}^{m \times n}$, where

| Ref. | Min./Max. | Guarantee | Complexity | D/R | Setting |
|---------|-------------|------------------|--|-----|----------|
| [13] | MAXAGREE | (E)PTAS | $n/\delta \cdot k^{O(\delta^{-2} \log k \log(1/\delta))}$ | R | k -CC |
| [13] | MINDISAGREE | PTAS | $n^{O(100k/\delta^2)} \cdot \log n$ | R | k -CC |
| [15] | MINDISAGREE | PTAS | $n^{O(9k/\delta^2)} \cdot \log n$ | R | k -CC |
| [11, 9] | MINDISAGREE | $O(\log n)$ -OPT | LP | D | Inc. CC |
| [17] | MAXAGREE | 0.766-OPT | SDP | D | Inc. CC |
| [5] | MINDISAGREE | 11-OPT | LP | D | BCC |
| [2] | MINDISAGREE | 4-OPT | LP | D | BCC |
| [2] | MINDISAGREE | 4-OPT | $ E $ | R | BCC |
| Ours | MAXAGREE | (E)PTAS | $2^{O(k/\delta^2 \cdot \log \sqrt{k}/\delta)} \cdot (\delta^{-2} + k) \cdot n + T_{\text{svd}}(\delta^{-2})$ | R | k -BCC |
| | MAXAGREE | (E)PTAS | $2^{O(\delta^{-3} \cdot \log \delta^{-3})} \cdot O(\delta^{-2}) \cdot n + T_{\text{svd}}(\delta^{-2})$ | R | BCC |

Table 1: Summary of results on BCC and related problems. For each scheme, we indicate the problem setting, the objective (MAXAGREE/MINDISAGREE), the guarantees (c -OPT implies a multiplicative factor approximation), and its computational complexity (n denotes the total number of vertices, LP/SDP denotes the complexity of a linear/semidefinite program, and $T_{\text{svd}}(r)$ the time required to compute a rank- r truncated SVD of a $n \times n$ matrix). The D/R column indicates whether the scheme is deterministic or randomized.

B_{ij} is equal to the weight of edge (i, j) .

Consider a clustering \mathcal{C} of the vertices $U \cup V$ into at most k clusters C_1, \dots, C_k . Let $\mathbf{X} \in \{0, 1\}^{m \times k}$ be the cluster assignment matrix for the vertices of U associated with \mathcal{C} , that is, $X_{ij} = 1$ if and only if vertex $i \in U$ is assigned to cluster C_j . Similarly, let $\mathbf{Y} \in \{0, 1\}^{n \times k}$ be the cluster assignment for V .

Lemma 2.1. *For any instance $G = (U, V, E)$ of the k -BCC problem with bi-adjacency matrix $\mathbf{B} \in \{\pm 1\}^{|U| \times |V|}$, and for any clustering \mathcal{C} of $U \cup V$ into k clusters, the number of agreements achieved by \mathcal{C} is*

$$\text{AGREE}(\mathcal{C}) = \text{Tr}(\mathbf{X}^\top \mathbf{B} \mathbf{Y}) + |E^-|,$$

where $\mathbf{X} \in \{0, 1\}^{|U| \times k}$ and $\mathbf{Y} \in \{0, 1\}^{|V| \times k}$ are the cluster assignment matrices corresponding to \mathcal{C} .

Proof. Let \mathbf{B}^+ and $\mathbf{B}^- \in \{0, 1\}^{|U| \times |V|}$ be indicator matrices for E^+ and E^- . Then, $\mathbf{B} = \mathbf{B}^+ - \mathbf{B}^-$. Given a clustering $\mathcal{C} = \{C_j\}_{j=1}^k$, or equivalently the assignment matrices $\mathbf{X} \in \{0, 1\}^{|U| \times k}$ and $\mathbf{Y} \in \{0, 1\}^{|V| \times k}$, the number of pairs of similar vertices assigned to the same cluster is equal to $\text{Tr}(\mathbf{X}^\top \mathbf{B}^+ \mathbf{Y})$. Similarly, the number of pairs of dissimilar vertices assigned to different clusters is equal to $|E^-| - \text{Tr}(\mathbf{X}^\top \mathbf{B}^- \mathbf{Y})$. The total number of agreements achieved by \mathcal{C} is

$$\begin{aligned} \text{AGREE}(\mathcal{C}) &= \text{Tr}(\mathbf{X}^\top \mathbf{B}^+ \mathbf{Y}) + |E^-| - \text{Tr}(\mathbf{X}^\top \mathbf{B}^- \mathbf{Y}) \\ &= \text{Tr}(\mathbf{X}^\top \mathbf{B} \mathbf{Y}) + |E^-|, \end{aligned}$$

which is the desired result. \square

It follows that computing a k -clustering that achieves the maximum number of agreements, boils down to a

constrained bilinear maximization:

$$\text{MAXAGREE}[k] = \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \mathbf{B} \mathbf{Y}) + |E^-|, \quad (2)$$

where

$$\begin{aligned} \mathcal{X} &\triangleq \{\mathbf{X} \in \{0, 1\}^{m \times k} : \|\mathbf{X}\|_{\infty, 1} = 1\}, \\ \mathcal{Y} &\triangleq \{\mathbf{Y} \in \{0, 1\}^{n \times k} : \|\mathbf{Y}\|_{\infty, 1} = 1\}. \end{aligned} \quad (3)$$

Here, $\|\mathbf{X}\|_{\infty, 1}$ denotes the maximum of the ℓ_1 -norm of the rows of \mathbf{X} . Since $\mathbf{X} \in \{0, 1\}^{|U| \times k}$, the constraint ensures that each row of \mathbf{X} has exactly one nonzero entry. Hence, \mathcal{X} and \mathcal{Y} describe the sets of valid cluster assignment matrices for U and V , respectively.

In the sequel, we briefly describe our approach for approximately solving the maximization in (2) under more general constraints, and subsequently apply it to the k -BCC/MAXAGREE problem.

3 Bilinear Maximization Framework

We describe a simple algorithm for computing an approximate solution to the constrained maximization

$$\max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y}), \quad (4)$$

where the input argument \mathbf{A} is a real $m \times n$ matrix, and \mathcal{X}, \mathcal{Y} are norm-bounded sets. Our approach is exponential in the rank of the argument \mathbf{A} , which can be prohibitive in practice. To mitigate this effect, we solve the maximization on a low-rank approximation $\tilde{\mathbf{A}}$ of \mathbf{A} , instead. The quality of the output depends on the spectrum of \mathbf{A} and the rank r of the surrogate matrix. Alg. 1 outlines our approach for solving (4), operating directly on the rank- r matrix $\tilde{\mathbf{A}}$.

If we knew the optimal value for variable \mathbf{Y} , then the optimal value for variable \mathbf{X} would be the solution to

$$P_{\mathcal{X}}(\mathbf{L}) = \arg \max_{\mathbf{X} \in \mathcal{X}} \text{TR}(\mathbf{X}^\top \mathbf{L}), \quad (5)$$

for $\mathbf{L} = \tilde{\mathbf{A}}\mathbf{Y}$. Similarly, with $\mathbf{R} = \tilde{\mathbf{A}}^\top \mathbf{X}$ for a given \mathbf{X} ,

$$P_{\mathcal{Y}}(\mathbf{R}) = \arg \max_{\mathbf{Y} \in \mathcal{Y}} \text{TR}(\mathbf{R}^\top \mathbf{Y}) \quad (6)$$

is the optimal value of \mathbf{Y} for that \mathbf{X} . Our algorithm requires that such a linear maximization oracles $P_{\mathcal{X}}(\cdot)$ and $P_{\mathcal{Y}}(\cdot)$ exist.² Of course, the optimal value for either of the two variables is not known. It is known, however, that the columns of the $m \times k$ matrix $\mathbf{L} = \tilde{\mathbf{A}}\mathbf{Y}$ lie in the r -dimensional range of $\tilde{\mathbf{A}}$ for all feasible $\mathbf{Y} \in \mathcal{Y}$. Alg. 1 effectively operates by sampling candidate values for \mathbf{L} . More specifically, it considers a large –exponential in $r \cdot k$ – collection of $r \times k$ matrices \mathbf{C} whose columns are points of an ϵ -net of the unit ℓ_2 -ball \mathbb{B}_2^{r-1} . Each matrix \mathbf{C} is used to generate a matrix \mathbf{L} whose k columns lie in the range of the input matrix $\tilde{\mathbf{A}}$ and in turn to produce a feasible solution pair \mathbf{X}, \mathbf{Y} by successively solving (5) and (6). Due to the properties of the ϵ -net, one of the computed feasible pairs is guaranteed to achieve an objective value in (4) close to the optimal one (for argument $\tilde{\mathbf{A}}$).

Lemma 3.2. *For any real $m \times n$, rank- r matrix $\tilde{\mathbf{A}}$, and sets $\mathcal{X} \subset \mathbb{R}^{m \times k}$ and $\mathcal{Y} \subset \mathbb{R}^{n \times k}$, let*

$$(\tilde{\mathbf{X}}_*, \tilde{\mathbf{Y}}_*) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{TR}(\mathbf{X}^\top \tilde{\mathbf{A}}\mathbf{Y}).$$

If there exist linear maximization oracles $P_{\mathcal{X}}$ and $P_{\mathcal{Y}}$ as in (5) and (6), then Alg. 1 with input $\tilde{\mathbf{A}}$ and accuracy $\epsilon \in (0, 1)$ outputs $\tilde{\mathbf{X}} \in \mathcal{X}$ and $\tilde{\mathbf{Y}} \in \mathcal{Y}$ such that

$$\text{TR}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}}\tilde{\mathbf{Y}}) \geq \text{TR}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{A}}\tilde{\mathbf{Y}}_*) - 2\epsilon\sqrt{k} \cdot \|\tilde{\mathbf{A}}\|_2 \cdot \mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}},$$

where $\mu_{\mathcal{X}} \triangleq \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_F$ and $\mu_{\mathcal{Y}} \triangleq \max_{\mathbf{Y} \in \mathcal{Y}} \|\mathbf{Y}\|_F$, in time $O((2\sqrt{r}/\epsilon)^{r \cdot k} \cdot (T_{\mathcal{X}} + T_{\mathcal{Y}} + (m+n)r)) + T_{\text{SVD}}(r)$. Here, $T_{\text{SVD}}(r)$ denotes the time required to compute the truncated SVD of $\tilde{\mathbf{A}}$, while $T_{\mathcal{X}}$ and $T_{\mathcal{Y}}$ denote the running time of the two oracles.

The crux of Lemma 3.2 is that if there exists an efficient procedure to solve the simpler maximizations (5) and (6), then we can approximately solve the bilinear maximization (4) in time that depends exponentially on the intrinsic dimension r of the input $\tilde{\mathbf{A}}$, but polynomially on the size of the input. A formal proof is given in Appendix Sec. A.

Recall that $\tilde{\mathbf{A}}$ is only a low-rank approximation of the potentially full rank original argument \mathbf{A} . The guarantees of Lemma 3.2 can be translated to guarantees

²In the case of the k -BCC problem (2), where \mathcal{X} and \mathcal{Y} correspond to the sets of cluster assignment matrices defined in (3), such optimization oracles exist (see Alg. 2).

Algorithm 1 BiLinearLowRankSolver

input : $m \times n$ real matrix $\tilde{\mathbf{A}}$ of rank r , $\epsilon \in (0, 1)$
output $\tilde{\mathbf{X}} \in \mathcal{X}$, $\tilde{\mathbf{Y}} \in \mathcal{Y}$ {See Lemma 3.2.}
 1: $\mathcal{C} \leftarrow \{\}$ {Candidate solutions}
 2: $\tilde{\mathbf{U}}, \tilde{\Sigma}, \tilde{\mathbf{V}} \leftarrow \text{SVD}(\tilde{\mathbf{A}})$ { $\tilde{\Sigma} \in \mathbb{R}^{r \times r}$ }
 3: **for each** $\mathbf{C} \in (\epsilon\text{-net of } \mathbb{B}_2^{r-1})^{\otimes k}$ **do**
 4: $\mathbf{L} \leftarrow \tilde{\mathbf{U}}\tilde{\Sigma}\mathbf{C}$ { $\mathbf{L} \in \mathbb{R}^{m \times k}$ }
 5: $\mathbf{X} \leftarrow P_{\mathcal{X}}(\mathbf{L})$
 6: $\mathbf{R} \leftarrow \mathbf{X}^\top \tilde{\mathbf{A}}$ { $\mathbf{R} \in \mathbb{R}^{k \times n}$ }
 7: $\mathbf{Y} \leftarrow P_{\mathcal{Y}}(\mathbf{R})$
 8: $\mathcal{C} \leftarrow \mathcal{C} \cup \{(\mathbf{X}, \mathbf{Y})\}$
 9: **end for**
 10: $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) \leftarrow \arg \max_{(\mathbf{X}, \mathbf{Y}) \in \mathcal{C}} \text{TR}(\mathbf{X}^\top \tilde{\mathbf{U}}\tilde{\Sigma}\tilde{\mathbf{V}}^\top \mathbf{Y})$

Algorithm 2 $P_{\mathcal{X}}(\cdot)$, $\mathcal{X} \triangleq \{\mathbf{X} \in \{0, 1\}^{m \times k} : \|\mathbf{X}\|_{\infty, 1} = 1\}$

input : $m \times k$ real matrix \mathbf{L}
output $\tilde{\mathbf{X}} = \arg \max_{\mathbf{X} \in \mathcal{X}} \text{TR}(\mathbf{X}^\top \mathbf{L})$
 1: $\tilde{\mathbf{X}} \leftarrow \mathbf{0}_{m \times k}$
 2: **for** $i = 1, \dots, m$ **do**
 3: $j_i \leftarrow \arg \max_{j \in [k]} L_{ij}$
 4: $\tilde{X}_{ij_i} \leftarrow 1$
 5: **end for**

for the original matrix introducing an additional error term to account for the extra level of approximation.

Lemma 3.3. *For any $\mathbf{A} \in \mathbb{R}^{m \times n}$, let*

$$(\mathbf{X}_*, \mathbf{Y}_*) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{TR}(\mathbf{X}^\top \mathbf{A}\mathbf{Y}),$$

where \mathcal{X} and \mathcal{Y} satisfy the conditions of Lemma 3.2. Let $\tilde{\mathbf{A}}$ be a rank- r approximation of \mathbf{A} , and $\tilde{\mathbf{X}} \in \mathcal{X}$, $\tilde{\mathbf{Y}} \in \mathcal{Y}$ be the output of Alg. 1 with input $\tilde{\mathbf{A}}$ and accuracy ϵ . Then,

$$\begin{aligned} & \text{TR}(\mathbf{X}_*^\top \mathbf{A}\mathbf{Y}_*) - \text{TR}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}}\tilde{\mathbf{Y}}) \\ & \leq 2 \cdot \left(\epsilon\sqrt{k} \cdot \|\tilde{\mathbf{A}}\|_2 + \|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \right) \cdot \mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}}. \end{aligned}$$

Lemma 3.3 follows from Lemma 3.2. A formal proof is deferred to Appendix Sec. A. This concludes the brief discussion on our bilinear maximization framework.

4 Our k -BCC Algorithm

The k -BCC/MAXAGREE problem on a bipartite graph $G = (U, V, E)$ can be written as a constrained bilinear maximization (2) on the bi-adjacency matrix \mathbf{B} over the sets of valid cluster assignment matrices (3) for V and U . Adopting the framework of the previous section to approximately solve the maximization in (2) we obtain our k -BCC algorithm.

The missing ingredient is a pair of efficient procedures $P_{\mathcal{X}}(\cdot)$ and $P_{\mathcal{Y}}(\cdot)$, as described in Lemma 3.2 for solving (5) and (6), respectively, in the case where \mathcal{X} and

\mathcal{Y} are the sets of valid cluster assignment matrices (3). Such a procedure exists and is outlined in Alg. 2.

Lemma 4.4. *For any $\mathbf{L} \in \mathbb{R}^{m \times k}$, Algorithm 2 outputs*

$$\arg \max_{\mathbf{X} \in \mathcal{X}} \text{Tr}(\mathbf{X}^\top \mathbf{L}),$$

where \mathcal{X} is defined in (3), in time $O(k \cdot m)$.

A proof is provided in Appendix, Sec. B.

Note that in our case, Alg. 2 is used as both $P_{\mathcal{X}}(\cdot)$ and $P_{\mathcal{Y}}(\cdot)$. Putting the pieces together, we obtain the core of our k -BCC algorithm, outlined in Alg. 3. Given a bipartite graph G , with bi-adjacency matrix \mathbf{B} , we first compute a rank- r approximation $\tilde{\mathbf{B}}$ of \mathbf{B} via the truncated SVD. Using Alg. 1, equipped with Alg. 2 as a subroutine, we approximately solve the bilinear maximization (2) with argument $\tilde{\mathbf{B}}$. The output is a pair of valid cluster assignment matrices for U and V .

Alg. 3 exposes the configuration parameters r and ϵ that control the performance of our bilinear solver, and in turn the quality of the output k -clustering. To simplify the description, we also create a k -BCC “wrapper” procedure outlined in Alg. 4: given a bound k on the number of clusters and a single accuracy parameter $\delta \in (0, 1)$, Alg. 4 configures and invokes Alg. 3.

Theorem 1. (k -BCC.) *For any instance $(G = (U, V, E), k)$ of the k -BCC problem with bi-adjacency matrix \mathbf{B} and for any desired accuracy parameter $\delta \in (0, 1)$, Algorithm 4 computes a clustering $\tilde{\mathcal{C}}^{(k)}$ of $U \cup V$ into at most k clusters, such that*

$$\text{AGREE}(\tilde{\mathcal{C}}^{(k)}) \geq (1 - \delta) \cdot \text{AGREE}(\mathcal{C}_\star^{(k)}),$$

where $\mathcal{C}_\star^{(k)}$ is the optimal k -clustering, in time $2^{O(k/\delta^2 \cdot \log \sqrt{k}/\delta)} \cdot (\delta^{-2} + k) \cdot (|U| + |V|) + T_{\text{SVD}}(\delta^{-2})$.

In the remainder of this section, we prove Theorem 1. We begin with the core Alg. 3, which invokes the low-rank bilinear solver Alg. 1 with accuracy parameters ϵ and r on the rank- r matrix $\tilde{\mathbf{B}}$ and computes a pair of valid cluster assignment matrices $\tilde{\mathbf{X}}^{(k)}$, $\tilde{\mathbf{Y}}^{(k)}$. By Lemma 3.2, and taking into account that $\sigma_1(\tilde{\mathbf{B}}) = \sigma_1(\mathbf{B}) \leq \|\mathbf{B}\|_F \leq \sqrt{mn}$, and the fact that $\|\mathbf{X}\|_F = \sqrt{m}$ and $\|\mathbf{Y}\|_F = \sqrt{n}$ for all valid cluster assignment matrix pairs, the output pair satisfies

$$\text{Tr}(\tilde{\mathbf{X}}^{(k)\top} \tilde{\mathbf{B}} \tilde{\mathbf{Y}}^{(k)}) - \text{Tr}(\tilde{\mathbf{X}}^{(k)\top} \tilde{\mathbf{B}} \tilde{\mathbf{Y}}^{(k)}) \leq 2 \cdot \epsilon \cdot \sqrt{k} \cdot mn,$$

where $(\tilde{\mathbf{X}}_\star^{(k)}, \tilde{\mathbf{Y}}_\star^{(k)}) \triangleq \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \tilde{\mathbf{B}} \mathbf{Y})$.

In turn, by Lemma 3.3 taking into account that $\|\mathbf{B} - \tilde{\mathbf{B}}\|_2 \leq \sqrt{mn}/\sqrt{r}$ (see Cor. 1), on the original bi-adjacency matrix \mathbf{B} the output pair satisfies

$$\begin{aligned} & \text{Tr}(\mathbf{X}_\star^{(k)\top} \mathbf{B} \mathbf{Y}_\star^{(k)}) - \text{Tr}(\tilde{\mathbf{X}}^{(k)\top} \tilde{\mathbf{B}} \tilde{\mathbf{Y}}^{(k)}) \\ & \leq 2\epsilon\sqrt{k} \cdot mn + 2(r+1)^{-1/2} \cdot mn, \end{aligned} \quad (7)$$

where $(\mathbf{X}_\star^{(k)}, \mathbf{Y}_\star^{(k)}) \triangleq \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \mathbf{B} \mathbf{Y})$. The unknown optimal pair $\mathbf{X}_\star^{(k)}, \mathbf{Y}_\star^{(k)}$ represents the optimal k -clustering $\mathcal{C}_\star^{(k)}$, *i.e.*, the clustering that achieves the maximum number of agreements using at most k clusters. Similarly, let $\tilde{\mathcal{C}}^{(k)}$ be the clustering induced by the computed pair $\tilde{\mathbf{X}}^{(k)}, \tilde{\mathbf{Y}}^{(k)}$. From (7) and Lemma 2.1, it immediately follows that

$$\begin{aligned} & \text{AGREE}(\mathcal{C}_\star^{(k)}) - \text{AGREE}(\tilde{\mathcal{C}}^{(k)}) \\ & \leq 2 \cdot (\epsilon\sqrt{k} + \sqrt{r+1}) \cdot mn. \end{aligned} \quad (8)$$

Eq. 8 establishes the guarantee of Alg. 3 as a function of its accuracy parameters r and ϵ . Substituting those parameters by the values assigned by Alg. 4, we obtain

$$\text{AGREE}(\mathcal{C}_\star^{(k)}) - \text{AGREE}(\tilde{\mathcal{C}}^{(k)}) \leq 2^{-1} \cdot \delta \cdot mn. \quad (9)$$

Fact 1. *For any BCC instance (or k -BCC with $k \geq 2$), the optimal clustering achieves at least $mn/2$ agreements.*

Proof. If more than half of the edges are labeled ‘+’, a single cluster containing all vertices achieves at least $nm/2$ agreements. Otherwise, two clusters corresponding to U and V achieve the same result. \square

In other words, Fact 1 states that

$$\text{AGREE}(\mathcal{C}_\star^{(k)}) \geq mn/2, \quad \forall k \geq 2. \quad (10)$$

Combining (10) with (9), we obtain

$$\text{AGREE}(\tilde{\mathcal{C}}^{(k)}) \geq (1 - \delta) \cdot \text{AGREE}(\mathcal{C}_\star^{(k)}),$$

which is the desired result. Finally, the computational complexity is obtained substituting the appropriate values in that of Alg. 1 as in Lemma 3.2, and the time complexity of subroutine Alg. 2 given in Lemma 4.4. This completes the proof of Theorem 1.

5 An Efficient PTAS for BCC

We provide an efficient polynomial time approximation scheme (EPTAS) for BCC/MAXAGREE, *i.e.*, the unconstrained version with no upper bound on the number of output clusters. We rely on the following key observation: *any constant factor approximation of the MAXAGREE objective, can be achieved by constant number of clusters.*³ Hence, the desired approximation algorithm for BCC can be obtained by invoking the k -BCC algorithm under the appropriate configuration.

Recall that in the unconstrained BCC setting, the number of output clusters is an optimization variable; the optimal clustering may comprise any number of clusters, from a single cluster containing all vertices of

³A similar result for CC/MAXAGREE exists in [7].

Algorithm 3 k -BCC via Low-rank Bilinear Maximization

input : • Bi-adjacency matrix $\mathbf{B} \in \{\pm 1\}^{m,n}$ of bipartite graph $G = (U, V, E)$,
 • Target number of clusters k , • Approximation rank r , • Accuracy parameter $\epsilon \in (0, 1)$.
output Clustering $\tilde{\mathcal{C}}^{(k)}$ of $U \cup V$ into k clusters –
 Equivalently, cluster membership matrices $\tilde{\mathbf{X}}^{(k)} \in \mathcal{X}$ and $\tilde{\mathbf{Y}}^{(k)} \in \mathcal{Y}$ for U and V , respectively.
 1: $\tilde{\mathbf{B}} \leftarrow \text{SVD}(\mathbf{B}, r)$ {Truncated SVD.}
 2: Approx. solve the bilinear maximization (2) on $\tilde{\mathbf{B}}$, over the sets \mathcal{X}, \mathcal{Y} of valid cluster assignment matrices (Eq. 3):
 $\tilde{\mathbf{X}}^{(k)}, \tilde{\mathbf{Y}}^{(k)} \leftarrow \text{BiLinearLowRankSolver}(\tilde{\mathbf{B}}, \epsilon, r)$ {Alg. 1, using Alg. 2 as $\text{P}_{\mathcal{X}}(\cdot)$ and $\text{P}_{\mathcal{Y}}(\cdot)$.}

Algorithm 4 k -BCC/MAXAGREE

input : Bi-adjacency matrix $\mathbf{B} \in \{\pm 1\}^{m,n}$,
 Target number of k ,
 Accuracy $\delta \in (0, 1)$.
output Clustering $\tilde{\mathcal{C}}^{(k)}$ of $U \cup V$ such that
 $\text{AGREE}(\tilde{\mathcal{C}}^{(k)}) \geq (1 - \delta) \cdot \text{AGREE}(\mathcal{C}_*^{(k)})$.
 1: Set up parameters:
 $\epsilon \leftarrow 2^{-3} \cdot \delta \cdot k^{-1/2}$, $r \leftarrow 2^6 \cdot \delta^{-2} - 1$.
 2: Return output of Alg. 3 for input $(\mathbf{B}, k, r, \epsilon)$.

Algorithm 5 A PTAS for BCC/MAXAGREE

input : Bi-adjacency matrix $\mathbf{B} \in \{\pm 1\}^{m,n}$,
 Accuracy $\delta \in (0, 1)$.
output Clustering $\tilde{\mathcal{C}}$ of $U \cup V$ (into at most
 $2^3 \cdot \delta^{-1}$ clusters) such that
 $\text{AGREE}(\tilde{\mathcal{C}}) \geq (1 - \delta) \cdot \text{AGREE}(\mathcal{C}_*)$.
 1: Set up parameters:
 $k \leftarrow 2^3 \cdot \delta^{-1}$, $\epsilon \leftarrow 2^{-6} \cdot \delta^2$, $r \leftarrow 2^8 \cdot \delta^{-2} - 1$.
 2: Return output of Alg. 3 for input $(\mathbf{B}, k, r, \epsilon)$.

the graph, up to $|U| + |V|$ singleton clusters. In principle, one could solve $\text{MAXAGREE}[k]$ for all possible values of the parameter k to identify the best clustering, but this approach is computationally intractable.

On the contrary, if one is willing to settle for an approximately optimal solution, then a constant number of clusters may suffice. More formally,

Lemma 5.5. *For any BCC instance, and $0 < \epsilon \leq 1$, there exists a clustering \mathcal{C} with at most $k = 2 \cdot \epsilon^{-1} + 2$ clusters such that $\text{AGREE}(\mathcal{C}) \geq \text{AGREE}(\mathcal{C}_*) - \epsilon \cdot nm$, where \mathcal{C}_* denotes the optimal clustering.*

We defer the proof to the end of the section.

In conjunction with Fact 1, Lemma 5.5 suggests that to obtain a constant factor approximation for the unconstrained BCC/MAXAGREE problem, for any constant arbitrarily close to 1, it suffices to solve a k -BCC instance for a sufficiently large –but constant– number of clusters k . In particular, to obtain an $(1 - \delta)$ -factor approximation for any constant $\delta \in (0, 1)$, it suffices to solve the k -BCC problem with an upper bound $k = O(\delta^{-1})$ on the number of clusters.

Alg. 5 outlines the approximation scheme for BCC. For a given accuracy parameter δ , it invokes Alg. 3 under an appropriate configuration of the parameters k, r and ϵ , yielding the following guarantees:

Theorem 2. (PTAS for BCC.) *For any instance $G = (U, V, E)$ of the BCC problem with bi-adjacency matrix \mathbf{B} and for any desired accuracy parameter $\delta \in (0, 1)$, Algorithm 5 computes a clustering $\tilde{\mathcal{C}}$ of $U \cup V$ into (at most) $2^3 \cdot \delta^{-1}$ clusters, such that*

$$\text{AGREE}(\tilde{\mathcal{C}}) \geq (1 - \delta) \cdot \text{AGREE}(\mathcal{C}_*),$$

where \mathcal{C}_* is an optimal clustering (with no constraint on the number of clusters), in time

$$2^{O(\delta^{-3} \cdot \log \delta^{-3})} \cdot \delta^{-2} \cdot (m + n) + T_{\text{svd}}(\delta^{-2}).$$

The proof of Theorem 2 follows from the guarantees of Alg. 3 in (8) substituting the values of the parameters k, r and ϵ with the values specified by Alg 5. The core k -BCC Alg. 3 returns a clustering $\tilde{\mathcal{C}}$ with at most $k_0 = 2^3 \cdot \delta^{-1}$ clusters that satisfies

$$\text{AGREE}(\tilde{\mathcal{C}}) \geq \text{AGREE}(\mathcal{C}_*^{(k_0)}) - \frac{\delta}{4} \cdot mn, \quad (11)$$

where $\mathcal{C}_*^{(k_0)}$ is the best among the clusterings using at most k_0 clusters. Also, for $k = k_0$, Lemma 5.5 implies that $\text{AGREE}(\mathcal{C}_*^{(k_0)}) \geq \text{AGREE}(\mathcal{C}_*) - \delta/4 \cdot nm$, where \mathcal{C}_* is the optimal clustering without any constraint on the number of output clusters. Hence,

$$\text{AGREE}(\tilde{\mathcal{C}}) \geq \text{AGREE}(\mathcal{C}_*) - \delta/2 \cdot mn. \quad (12)$$

Continuing from (12), and taking into account Fact 1, we conclude that the output $\tilde{\mathcal{C}}$ of Alg. 5 satisfies

$$\text{AGREE}(\tilde{\mathcal{C}}) \geq (1 - \delta) \cdot \text{AGREE}(\mathcal{C}_*), \quad (13)$$

which is the desired result. Finally, the computational complexity of Alg. 5 follows from that of Alg. 3 substituting the parameter values in Lemma 3.2, taking into account the running time $T_{\mathcal{X}} = O(\delta^{-2} \cdot m)$ of Alg. 2 used as subroutine $\text{P}_{\mathcal{X}}(\cdot)$ and similarly $T_{\mathcal{Y}} = O(\delta^{-2} \cdot n)$ for $\text{P}_{\mathcal{Y}}(\cdot)$. This concludes the proof of Theorem 2.

For any desired constant accuracy $\delta \in (0, 1)$, Alg. 5 outputs a clustering that achieves a number of agreements within a $(1 - \delta)$ -factor from the optimal, in time

that grows exponentially in δ^{-1} , but linearly in the size mn of the input. In other words, Alg. 5 is an EPTAS for BCC/MAXAGREE.

We complete the section with a proof for Lemma 5.5, which stated that a constant number of clusters suffices to obtain an approximately optimal solution.

5.1 Proof of Lemma 5.5

It suffices to consider $\epsilon > 2/(m + n - 2)$ as the lemma holds trivially otherwise. Without loss of generality, we focus on clusterings whose all but at most two clusters contain vertices from both U and V ; one of the two remaining clusters can contain vertices only from U and the other only from V . To verify that, consider an arbitrary clustering \mathcal{C} and let \mathcal{C}' be the clustering obtained by merging all clusters of \mathcal{C} containing only vertices of U into a single cluster, and those containing only vertices of V into another. The number of agreements is not affected, *i.e.*, $\text{AGREE}(\mathcal{C}') = \text{AGREE}(\mathcal{C})$.

We show that a clustering \mathcal{C} with the properties described in the lemma exists by modifying \mathcal{C}_* . Let $\mathcal{S}_U \subseteq \mathcal{C}_*$ be the set of clusters that contain at most $\epsilon m/2$ vertices of U , and $\mathcal{S}_V \subseteq \mathcal{C}_*$ those containing at most $\epsilon n/2$ vertices of V . (\mathcal{S}_U and \mathcal{S}_V may not be disjoint). Finally, let \mathcal{B} be the set of vertices contained in clusters of $\mathcal{S}_U \cup \mathcal{S}_V$. We construct \mathcal{C} as follows. Clusters of \mathcal{C}_* not in $\mathcal{S}_U \cup \mathcal{S}_V$ are left intact. Each such cluster has size at least $\epsilon(n + m)/2$. Further, all vertices in \mathcal{B} are rearranged into two clusters: one for the vertices in $\mathcal{B} \cap U$ and one those in $\mathcal{B} \cap V$.

The above rearrangement can reduce the number of agreements by at most ϵnm . To verify that, consider a cluster C in \mathcal{S}_U ; the cluster contains at most $\epsilon m/2$ vertices of U . Let $t \triangleq |V \cap C|$. Splitting C into two smaller clusters $C \cap U$ and $C \cap V$ can reduce the number of agreements by $\frac{1}{2}\epsilon mt$, *i.e.*, the total number of edges in C . Note that agreements on $-$ edges are not affected by splitting a cluster. Performing the same operation on all clusters in \mathcal{S}_U incurs a total reduction

$$\frac{1}{2}\epsilon m \sum_{C \in \mathcal{S}_U} |V \cap C| \leq \frac{1}{2}\epsilon m |V| = \frac{1}{2}\epsilon mn.$$

Repeating on \mathcal{S}_V incurs an additional cost of at most $\epsilon mn/2$ agreements, while merging all clusters containing only vertices from U (similarly for V) into a single cluster does not reduce the agreements. We conclude that $\text{AGREE}(\mathcal{C}) \geq \text{AGREE}(\mathcal{C}_*) - \epsilon nm$.

Finally, by construction all but at most two clusters in \mathcal{C} contain at least $\frac{1}{2}\epsilon(n + m)$. In turn,

$$n + m = \sum_{C \in \mathcal{C}} |C| \geq (|\mathcal{C}| - 2) \cdot \frac{1}{2}\epsilon(n + m),$$

from which the desired result follows. \square

6 Experiments

We evaluate our algorithm on synthetic and real data and compare with PivotBiCluster [2] and the SDP-based approach of [17]⁴ Although [5, 9, 11] provide algorithms applicable to BCC, those rely on LP or SDP formulations with a high number of constraints which imposes limitations in practice.

We run our core k -BCC algorithm (Alg. 3) to obtain a k -clustering. Disregarding the theoretical guarantees, we apply a threshold on the execution time which can only hurt the performance; equivalently, the iterative procedure of Alg. 1 is executed for an arbitrary subset of the points in the ϵ -net (Alg. 1, step 3).

6.1 Synthetic Data

We generate synthetic BCC instances as follows. We arbitrarily set $m = 100$, $n = 50$, and number of clusters $k' = 5$, and construct a complete bipartite graph $G = (U, V, E)$ with $|U| = m$, $|V| = n$ and assign binary ± 1 labels to the edges according to a random k' -clustering of the vertices. Then, we modify the sign of each label independently with probability p . We consider multiple values of p in the range $[0, 0.5]$. For each value, we generate 10 random instances as described above and compute a vertex clustering using all algorithms.

PivotBiCluster returns a clustering of arbitrary size; no parameter k is specified and the algorithm determines the number of clusters in an attempt to maximize the number of agreements. The majority of the output clusters are singletons which can be merged into two clusters (see Subsec. 5.1). The algorithm is substantially faster than the other approaches on examples of this scale. Hence, we run PivotBiCluster with $3 \cdot 10^4$ random restarts on each instance and depict best results to allow for comparable running times. Contrary to PivotBiCluster, for our algorithm, which is referred to as BccBilinear, we need to specify the target number k of clusters. We run it for $k = 5, 10$ and 15 and arbitrarily set the parameter $r = 5$ and terminate our algorithm after 10^4 random samples/rounds. Finally, the SDP-based approach returns 4 clusters by construction, and is configured to output best results among 100 random pivoting steps.

Fig. 1 depicts the number of agreements achieved by each algorithm for each value of p , the number of output clusters, and the execution times. All numbers are averages over multiple random instances.

We note that in some cases and contrary to intuition,

⁴All algorithms are prototypically implemented in Matlab. [17] was implemented using CVX.




Figure 1: Synthetic data. We generate an $m \times n$ bipartite graph with edge weights ± 1 according to a random vertex k -clustering and subsequently flip the sign of each edge with probability p . We plot the average number of agreements achieved by each method over 10 random instances for each p value. The bar plot depicts the average number of output clusters for each scheme/ p -value pair. Within each bar, a horizontal line marks the number of singleton clusters.

our algorithm performs better for lower values of k , the target number of clusters. We attribute this phenomenon to the fact that for higher values of k , we should typically use higher approximation rank and consider larger number of samples.

6.2 MovieLens Dataset

The MovieLens datasets [14] are sets of movie ratings: each of m users assigns scores in $\{1, \dots, 5\}$ to a small subset of the n movies. Table 2 lists the dimensions of the datasets. From each dataset, we generate an instance of the (incomplete) BCC problem as follows: we construct the bipartite graph on the user-movie pairs using the ratings as edge weights, compute the average weight and finally set weights higher than the average to $+1$ and the others to -1 .

We run our algorithm to obtain a clustering of the ver-

| Dataset | m (Users) | n (Movies) | Ratings |
|---------------|-------------|--------------|---------|
| MovieLens100K | 1000 | 1700 | 10^5 |
| MovieLens1M | 6000 | 4000 | 10^6 |
| MovieLens10M | 72000 | 10000 | 10^7 |

Table 2: Summary of MovieLens datasets [14].

ties. For a reference, we compare to PivotBiCluster with 50 random restarts. Note, however, that PivotBiCluster is not designed for the incomplete BCC problem; to apply the algorithm, we effectively treat missing edges as edges of negative weight. Finally, the SDP approach of [17], albeit suitable for the incomplete CC problem, does not scale to this size of input. Table 3 lists the number of agreements achieved by each method on each one of the three datasets and the corresponding execution times.

7 Conclusions

We presented the first algorithm with provable approximation guarantees for k -BCC/MAXAGREE. Our approach relied on formulating k -BCC as a constrained bilinear maximization over the sets of cluster assignment matrices and developing a simple framework to approximately solve that combinatorial optimization.

In the unconstrained BCC setting, with no bound on the number of output clusters, we showed that any constant multiplicative factor approximation for the MAXAGREE objective can be achieved using a constant number of clusters. In turn, under the appropriate configuration, our k -BCC algorithm yields an Efficient PTAS for BCC/MAXAGREE.

Acknowledgments This research has been supported by NSF Grants CCF 1344179, 1344364, 1407278, 1422549 and ARO YIP W911NF-14-1-0258. DP is supported by NSF awards CCF-1217058 and CCF-1116404 and MURI AFOSR grant 556016.

| MovieLens | 100K | 1M | 10M |
|----------------|------------------|--------------------|---------------------------------|
| PivotBiCluster | 46134 (27.95) | 429277 (651.13) | 5008577 ($1.5 \cdot 10^5$) |
| BccBilinear | 68141 (6.65) | 694366 (19.50) | 6857509 ($1.2 \cdot 10^3$) |

Table 3: Number of agreements achieved by the two algorithms on incomplete (k)-BCC instances obtained from the MovieLens datasets [14]. For PivotBiCluster we present best results over 50 random restarts. Our algorithm was arbitrarily configured with $r = 4$, $k = 10$ and a limit of 10^4 samples (iterations). We also note in parentheses the runtimes in seconds.

References

- [1] KookJin Ahn, Graham Cormode, Sudipto Guha, Andrew McGregor, and Anthony Wirth. Correlation clustering in data streams. In *Proceedings of the 32nd International Conference on Machine Learning (ICML-15)*, pages 2237–2246, 2015.
- [2] Nir Ailon, Noa Avigdor-Elgrabli, Edo Liberty, and Anke Van Zuylen. Improved approximation algorithms for bipartite correlation clustering. *SIAM Journal on Computing*, 41(5):1110–1121, 2012.
- [3] Nir Ailon, Moses Charikar, and Alantha Newman. Aggregating inconsistent information: ranking and clustering. *Journal of the ACM (JACM)*, 55(5):23, 2008.
- [4] Nir Ailon and Edo Liberty. Correlation clustering revisited: The “true” cost of error minimization problems. In *Automata, Languages and Programming*, pages 24–36. Springer, 2009.
- [5] Noga Amit. The bicluster graph editing problem. Master’s thesis, Tel Aviv University, 2004.
- [6] Megasthenis Asteris, Dimitris Papailiopoulos, Anastasios Kyrillidis, and Alexandros G Dimakis. Sparse pca via bipartite matchings. In *Advances in Neural Information Processing Systems 28*, pages 766–774. Curran Associates, Inc., 2015.
- [7] Nikhil Bansal, Avrim Blum, and Shuchi Chawla. Correlation clustering. *Machine Learning*, 56(1-3):89–113, 2004.
- [8] Francesco Bonchi, David Garcia-Soriano, and Edo Liberty. Correlation clustering: from theory to practice. In *Proceedings of the 20th ACM International Conference on Knowledge Discovery and Data mining*, pages 1972–1972. ACM, 2014.
- [9] Moses Charikar, Venkatesan Guruswami, and Anthony Wirth. Clustering with qualitative information. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 524–533. IEEE, 2003.
- [10] William W Cohen and Jacob Richman. Learning to match and cluster large high-dimensional data sets for data integration. In *Proceedings of the 8th ACM International Conference on Knowledge Discovery and Data mining*, pages 475–480. ACM, 2002.
- [11] Erik D Demaine, Dotan Emanuel, Amos Fiat, and Nicole Immorlica. Correlation clustering in general weighted graphs. *Theoretical Computer Science*, 361(2):172–187, 2006.
- [12] Xiaoli Zhang Fern and Carla E Brodley. Solving cluster ensemble problems by bipartite graph partitioning. In *Proceedings of the 21st International Conference on Machine Learning*, page 36. ACM, 2004.
- [13] Ioannis Giotis and Venkatesan Guruswami. Correlation clustering with a fixed number of clusters. In *Proceedings of the 17th annual ACM-SIAM Symposium on Discrete algorithms*, pages 1167–1176. Society for Industrial and Applied Mathematics, 2006.
- [14] University of Minnesota GroupLens Lab. MovieLens datasets. <http://grouplens.org/datasets/movielens/>, 2015. Accessed: 2015-10-03.
- [15] Marek Karpinski and Warren Schudy. Linear time approximation schemes for the gale-berlekamp game and related minimization problems. In *Proceedings of the 41st annual ACM Symposium on Theory of Computing*, pages 313–322. ACM, 2009.
- [16] Sara C Madeira and Arlindo L Oliveira. Biclustering algorithms for biological data analysis: a survey. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 1(1):24–45, 2004.
- [17] Chaitanya Swamy. Correlation clustering: maximizing agreements via semidefinite programming. In *Proceedings of the 15th annual ACM-SIAM Symposium on Discrete Algorithms*, pages 526–527. Society for Industrial and Applied Mathematics, 2004.
- [18] Panagiotis Symeonidis, Alexandros Nanopoulos, Apostolos Papadopoulos, and Yannis Manolopoulos. Nearest-biclusters collaborative filtering with constant values. In *Advances in web mining and web usage analysis*, pages 36–55. Springer, 2007.
- [19] Anke Van Zuylen and David P Williamson. Deterministic pivoting algorithms for constrained ranking and clustering problems. *Mathematics of Operations Research*, 34(3):594–620, 2009.
- [20] Michail Vlachos, Francesco Fusco, Charalambos Mavroforakis, Anastasios Kyrillidis, and Vassilios G Vassiliadis. Improving co-cluster quality with application to product recommendations. In *Proceedings of the 23rd ACM International Conference on Information and Knowledge Management*, pages 679–688. ACM, 2014.
- [21] Hongyuan Zha, Xiaofeng He, Chris Ding, Horst Simon, and Ming Gu. Bipartite graph partitioning and data clustering. In *Proceedings of the 10th ACM International Conference on Information and Knowledge Management*, pages 25–32. ACM, 2001.

A Bilinear Maximization Guarantees

Lemma 3.2. For any real $m \times n$, rank- r matrix $\tilde{\mathbf{A}}$ and arbitrary norm-bounded sets $\mathcal{X} \subset \mathbb{R}^{m \times k}$ and $\mathcal{Y} \subset \mathbb{R}^{n \times k}$, let

$$(\tilde{\mathbf{X}}_*, \tilde{\mathbf{Y}}_*) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \tilde{\mathbf{A}} \mathbf{Y}).$$

If there exist operators $P_{\mathcal{X}} : \mathbb{R}^{m \times k} \rightarrow \mathcal{X}$ such that

$$P_{\mathcal{X}}(\mathbf{L}) = \arg \max_{\mathbf{X} \in \mathcal{X}} \text{Tr}(\mathbf{X}^\top \mathbf{L})$$

and similarly, $P_{\mathcal{Y}} : \mathbb{R}^{n \times k} \rightarrow \mathcal{Y}$ such that

$$P_{\mathcal{Y}}(\mathbf{R}) = \arg \max_{\mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{R}^\top \mathbf{Y})$$

with running times $T_{\mathcal{X}}$ and $T_{\mathcal{Y}}$, respectively, then Algorithm 1 outputs $\tilde{\mathbf{X}} \in \mathcal{X}$ and $\tilde{\mathbf{Y}} \in \mathcal{Y}$ such that

$$\text{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) \geq \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_*) - 2\epsilon\sqrt{k} \cdot \|\tilde{\mathbf{A}}\|_2 \cdot \mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}},$$

where $\mu_{\mathcal{X}} \triangleq \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_{\text{F}}$ and $\mu_{\mathcal{Y}} \triangleq \max_{\mathbf{Y} \in \mathcal{Y}} \|\mathbf{Y}\|_{\text{F}}$, in time $O((2\sqrt{r}/\epsilon)^{r \cdot k} \cdot (T_{\mathcal{X}} + T_{\mathcal{Y}} + (m+n)r)) + T_{\text{svd}}(r)$.

Proof. In the sequel, $\tilde{\mathbf{U}}$, $\tilde{\Sigma}$ and $\tilde{\mathbf{V}}$ are used to denote the r -truncated singular value decomposition of $\tilde{\mathbf{A}}$.

Without loss of generality, we assume that $\mu_{\mathcal{X}} = \mu_{\mathcal{Y}} = 1$ since the variables in \mathcal{X} and \mathcal{Y} can be normalized by $\mu_{\mathcal{X}}$ and $\mu_{\mathcal{Y}}$, respectively, while simultaneously scaling the singular values of $\tilde{\mathbf{A}}$ by a factor of $\mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}}$. Then, $\|\mathbf{Y}\|_{\infty, 2} \leq 1$, $\forall \mathbf{Y} \in \mathcal{Y}$, where $\|\mathbf{Y}\|_{\infty, 2}$ denotes the maximum of the ℓ_2 -norm of the columns of \mathbf{Y} .

Let $\tilde{\mathbf{X}}_*, \tilde{\mathbf{Y}}_*$ be the optimal pair on $\tilde{\mathbf{A}}$, i.e.,

$$(\tilde{\mathbf{X}}_*, \tilde{\mathbf{Y}}_*) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}^\top \tilde{\mathbf{A}} \mathbf{Y})$$

and define the $r \times k$ matrix $\tilde{\mathbf{C}}_* \triangleq \tilde{\mathbf{V}}^\top \tilde{\mathbf{Y}}_*$. Note that

$$\begin{aligned} \|\tilde{\mathbf{C}}_*\|_{\infty, 2} &= \|\tilde{\mathbf{V}}^\top \tilde{\mathbf{Y}}_*\|_{\infty, 2} \\ &= \max_{1 \leq i \leq k} \|\tilde{\mathbf{V}}^\top [\tilde{\mathbf{Y}}_*]_{:, i}\|_2 \leq 1, \end{aligned} \quad (14)$$

with the last inequality following from the facts that $\|\mathbf{Y}\|_{\infty, 2} \leq 1 \forall \mathbf{Y} \in \mathcal{Y}$ and the columns of $\tilde{\mathbf{V}}$ are orthonormal. Alg. 1 iterates over the points in $(\mathbb{B}_2^{r-1})^{\otimes k}$. The latter is used to describe the set of $r \times k$ matrices whose columns have ℓ_2 norm at most equal to 1. At each point, the algorithm computes a candidate solution. By (14), the ϵ -net contains an $r \times k$ matrix $\mathbf{C}_\#$ such that

$$\|\mathbf{C}_\# - \tilde{\mathbf{C}}_*\|_{\infty, 2} \leq \epsilon.$$

Let $\mathbf{X}_\#, \mathbf{Y}_\#$ be the candidate pair computed at $\mathbf{C}_\#$ by the two step maximization, i.e.,

$$\mathbf{X}_\# \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}} \text{Tr}(\mathbf{X}^\top \tilde{\mathbf{U}} \tilde{\Sigma} \mathbf{C}_\#)$$

and

$$\mathbf{Y}_\# \triangleq \arg \max_{\mathbf{Y} \in \mathcal{Y}} \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{A}} \mathbf{Y}). \quad (15)$$

We show that the objective values achieved by the candidate pair $\mathbf{X}_\#, \mathbf{Y}_\#$ satisfies the inequality of the lemma implying the desired result.

By the definition of $\tilde{\mathbf{C}}_*$ and the linearity of the trace,

$$\begin{aligned} &\text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_*) \\ &= \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{C}}_*) \\ &= \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} \mathbf{C}_\#) + \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\tilde{\mathbf{C}}_* - \mathbf{C}_\#)) \\ &\leq \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} \mathbf{C}_\#) + \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\tilde{\mathbf{C}}_* - \mathbf{C}_\#)). \end{aligned} \quad (16)$$

The inequality follows from the fact that (by definition (15)) $\mathbf{X}_\#$ maximizes the first term over all $\mathbf{X} \in \mathcal{X}$. We compute an upper bound on the right hand side of (16). Define

$$\hat{\mathbf{Y}} \triangleq \arg \min_{\mathbf{Y} \in \mathcal{Y}} \|\tilde{\mathbf{V}}^\top \mathbf{Y} - \mathbf{C}_\#\|_{\infty, 2}.$$

(We note that $\hat{\mathbf{Y}}$ is used for the analysis and is never explicitly calculated.) Further, define the $r \times k$ matrix $\hat{\mathbf{C}} \triangleq \tilde{\mathbf{V}}^\top \hat{\mathbf{Y}}$. By the linearity of the trace operator

$$\begin{aligned} &\text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} \mathbf{C}_\#) \\ &= \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} \hat{\mathbf{C}}) + \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\mathbf{C}_\# - \hat{\mathbf{C}})) \\ &= \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top \hat{\mathbf{Y}}) + \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\mathbf{C}_\# - \hat{\mathbf{C}})) \\ &\leq \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top \mathbf{Y}_\#) + \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\mathbf{C}_\# - \hat{\mathbf{C}})) \\ &= \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{A}} \mathbf{Y}_\#) + \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\mathbf{C}_\# - \hat{\mathbf{C}})). \end{aligned} \quad (17)$$

The inequality follows from the fact that (by definition (15)) $\mathbf{Y}_\#$ maximizes the first term over all $\mathbf{Y} \in \mathcal{Y}$. Combining (17) and (16), and rearranging the terms, we obtain

$$\begin{aligned} &\text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_*) - \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{A}} \mathbf{Y}_\#) \\ &\leq \text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\tilde{\mathbf{C}}_* - \mathbf{C}_\#)) + \text{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\mathbf{C}_\# - \hat{\mathbf{C}})). \end{aligned} \quad (18)$$

By Lemma C.10,

$$\begin{aligned} &|\text{Tr}(\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\tilde{\mathbf{C}}_* - \mathbf{C}_\#))| \\ &\leq \|\tilde{\mathbf{X}}_*^\top \tilde{\mathbf{U}}\|_{\text{F}} \cdot \|\tilde{\Sigma}\|_2 \cdot \|\tilde{\mathbf{C}}_* - \mathbf{C}_\#\|_{\text{F}} \\ &\leq \|\tilde{\mathbf{X}}_*\|_{\text{F}} \cdot \sigma_1(\tilde{\mathbf{A}}) \cdot \sqrt{k} \cdot \epsilon \\ &\leq \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_{\text{F}} \cdot \sigma_1(\tilde{\mathbf{A}}) \cdot \sqrt{k} \cdot \epsilon \\ &\leq \sigma_1(\tilde{\mathbf{A}}) \cdot \sqrt{k} \cdot \epsilon. \end{aligned} \quad (19)$$

Similarly,

$$\begin{aligned} |\mathrm{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{U}} \tilde{\Sigma} (\mathbf{C}_\# - \hat{\mathbf{C}}))| &\leq \|\mathbf{X}_\# \tilde{\mathbf{U}}\|_F \cdot \|\tilde{\Sigma}\|_2 \cdot \|\mathbf{C}_\# - \hat{\mathbf{C}}\|_F \\ &\leq \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_F \cdot \sigma_1(\tilde{\mathbf{A}}) \cdot \sqrt{k} \cdot \epsilon \\ &\leq \sigma_1(\tilde{\mathbf{A}}) \cdot \sqrt{k} \cdot \epsilon. \end{aligned} \quad (20)$$

The second inequality follows from the fact that by the definition of $\hat{\mathbf{C}}$,

$$\begin{aligned} \|\hat{\mathbf{C}} - \mathbf{C}_\# \|_{\infty,2} &= \|\tilde{\mathbf{V}}^\top \hat{\mathbf{Y}} - \mathbf{C}_\# \|_{\infty,2} \leq \|\tilde{\mathbf{V}}^\top \tilde{\mathbf{Y}}_\star - \mathbf{C}_\# \|_{\infty,2} \\ &= \|\tilde{\mathbf{C}}_\star - \mathbf{C}_\# \|_{\infty,2} \leq \epsilon, \end{aligned}$$

which implies that

$$\|\hat{\mathbf{C}} - \mathbf{C}_\# \|_F \leq \sqrt{k} \cdot \epsilon.$$

Continuing from (18) under (19) and (20),

$$\mathrm{Tr}(\mathbf{X}_\#^\top \tilde{\mathbf{A}} \mathbf{Y}_\#) \geq \mathrm{Tr}(\tilde{\mathbf{X}}_\star^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_\star) - 2 \cdot \epsilon \cdot \sqrt{k} \cdot \sigma_1(\tilde{\mathbf{A}}).$$

Recalling that the singular values of $\tilde{\mathbf{A}}$ have been scaled by a factor of $\mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}}$ yields the desired result.

The runtime of Alg. 1 follows from the cost per iteration and the cardinality of the ϵ -net. Matrix multiplications can exploit the truncated singular value decomposition of $\tilde{\mathbf{A}}$ which is performed only once. \square

Lemma A.6. *For any $\mathbf{A}, \tilde{\mathbf{A}} \in \mathbb{R}^{m \times n}$, and norm-bounded sets $\mathcal{X} \subseteq \mathbb{R}^{m \times k}$ and $\mathcal{Y} \subseteq \mathbb{R}^{n \times k}$, let*

$$(\mathbf{X}_\star, \mathbf{Y}_\star) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \mathrm{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y}),$$

and

$$(\tilde{\mathbf{X}}_\star, \tilde{\mathbf{Y}}_\star) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \mathrm{Tr}(\mathbf{X}^\top \tilde{\mathbf{A}} \mathbf{Y}).$$

For any $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) \in \mathcal{X} \times \mathcal{Y}$ such that

$$\mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) \geq \gamma \cdot \mathrm{Tr}(\tilde{\mathbf{X}}_\star^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_\star) - C$$

for some $0 < \gamma \leq 1$, we have

$$\begin{aligned} \mathrm{Tr}(\tilde{\mathbf{X}}^\top \mathbf{A} \tilde{\mathbf{Y}}) &\geq \gamma \cdot \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - C \\ &\quad - 2 \cdot \|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \cdot \mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}}. \end{aligned}$$

where $\mu_{\mathcal{X}} \triangleq \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_F$ and $\mu_{\mathcal{Y}} \triangleq \max_{\mathbf{Y} \in \mathcal{Y}} \|\mathbf{Y}\|_F$.

Proof. By the optimality of $\tilde{\mathbf{X}}_\star, \tilde{\mathbf{Y}}_\star$ for $\tilde{\mathbf{A}}$, we have

$$\mathrm{Tr}(\tilde{\mathbf{X}}_\star^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_\star) \geq \mathrm{Tr}(\mathbf{X}_\star^\top \tilde{\mathbf{A}} \mathbf{Y}_\star).$$

In turn, for any $(\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}) \in \mathcal{X} \times \mathcal{Y}$ such that

$$\mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) \geq \gamma \cdot \mathrm{Tr}(\tilde{\mathbf{X}}_\star^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}_\star) - C$$

for some $0 < \gamma < 1$ (if such pairs exist), we have

$$\mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) \geq \gamma \cdot \mathrm{Tr}(\mathbf{X}_\star^\top \tilde{\mathbf{A}} \mathbf{Y}_\star) - C. \quad (21)$$

By the linearity of the trace operator,

$$\begin{aligned} \mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) &= \mathrm{Tr}(\tilde{\mathbf{X}}^\top \mathbf{A} \tilde{\mathbf{Y}}) - \mathrm{Tr}(\tilde{\mathbf{X}}^\top (\mathbf{A} - \tilde{\mathbf{A}}) \tilde{\mathbf{Y}}) \\ &\leq \mathrm{Tr}(\tilde{\mathbf{X}}^\top \mathbf{A} \tilde{\mathbf{Y}}) + |\mathrm{Tr}(\tilde{\mathbf{X}}^\top (\mathbf{A} - \tilde{\mathbf{A}}) \tilde{\mathbf{Y}})|. \end{aligned} \quad (22)$$

By Lemma C.10,

$$\begin{aligned} |\mathrm{Tr}(\tilde{\mathbf{X}}^\top (\mathbf{A} - \tilde{\mathbf{A}}) \tilde{\mathbf{Y}})| &\leq \|\tilde{\mathbf{X}}\|_F \cdot \|\tilde{\mathbf{Y}}\|_F \cdot \|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \\ &\leq \|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \cdot \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_F \cdot \max_{\mathbf{Y} \in \mathcal{Y}} \|\mathbf{Y}\|_F \triangleq R. \end{aligned} \quad (23)$$

Continuing from (22),

$$\mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) \leq \mathrm{Tr}(\tilde{\mathbf{X}}^\top \mathbf{A} \tilde{\mathbf{Y}}) + R. \quad (24)$$

Similarly,

$$\begin{aligned} \mathrm{Tr}(\mathbf{X}_\star^\top \tilde{\mathbf{A}} \mathbf{Y}_\star) &= \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - \mathrm{Tr}(\mathbf{X}_\star^\top (\mathbf{A} - \tilde{\mathbf{A}}) \mathbf{Y}_\star) \\ &\geq \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - |\mathrm{Tr}(\mathbf{X}_\star^\top (\mathbf{A} - \tilde{\mathbf{A}}) \mathbf{Y}_\star)| \\ &\geq \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - R. \end{aligned} \quad (25)$$

Combining the above, we have

$$\begin{aligned} \mathrm{Tr}(\tilde{\mathbf{X}}^\top \mathbf{A} \tilde{\mathbf{Y}}) &\geq \mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) - R \\ &\geq \gamma \cdot \mathrm{Tr}(\mathbf{X}_\star^\top \tilde{\mathbf{A}} \mathbf{Y}_\star) - R - C \\ &\geq \gamma \cdot (\mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - R) - R - C \\ &= \gamma \cdot \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - (1 + \gamma) \cdot R - C \\ &\geq \gamma \cdot \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - 2 \cdot R - C, \end{aligned}$$

where the first inequality follows from (24) the second from (21), the third from (25), and the last from the fact that $R \geq 0$. This concludes the proof. \square

Lemma 3.3. *For any $\mathbf{A} \in \mathbb{R}^{m \times n}$, let*

$$(\mathbf{X}_\star, \mathbf{Y}_\star) \triangleq \arg \max_{\mathbf{X} \in \mathcal{X}, \mathbf{Y} \in \mathcal{Y}} \mathrm{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y}),$$

where $\mathcal{X} \subseteq \mathbb{R}^{m \times k}$ and $\mathcal{Y} \subseteq \mathbb{R}^{n \times k}$ are sets satisfying the conditions of Lemma 3.2. Let $\tilde{\mathbf{A}}$ be a rank- r approximation of \mathbf{A} , and $\tilde{\mathbf{X}} \in \mathcal{X}$, $\tilde{\mathbf{Y}} \in \mathcal{Y}$ be the output of Alg. 1 with input $\tilde{\mathbf{A}}$ and accuracy ϵ . Then,

$$\begin{aligned} \mathrm{Tr}(\mathbf{X}_\star^\top \mathbf{A} \mathbf{Y}_\star) - \mathrm{Tr}(\tilde{\mathbf{X}}^\top \tilde{\mathbf{A}} \tilde{\mathbf{Y}}) &\leq 2 \cdot \left(\epsilon \sqrt{k} \cdot \|\tilde{\mathbf{A}}\|_2 + \|\mathbf{A} - \tilde{\mathbf{A}}\|_2 \right) \cdot \mu_{\mathcal{X}} \cdot \mu_{\mathcal{Y}}, \end{aligned}$$

where $\mu_{\mathcal{X}} \triangleq \max_{\mathbf{X} \in \mathcal{X}} \|\mathbf{X}\|_F$ and $\mu_{\mathcal{Y}} \triangleq \max_{\mathbf{Y} \in \mathcal{Y}} \|\mathbf{Y}\|_F$.

Proof. The proof follows the approximation guarantees of Alg. 1 in Lemma 3.2 and Lemma A.6. \square

B Correctness of Algorithm 2

In the sequel, we use $\|\mathbf{X}\|_{\infty,1}$ to denote the maximum of the ℓ_1 norm of the rows of \mathbf{X} . When $\mathbf{X} \in \{0,1\}^{d \times k}$, the constraint $\|\mathbf{X}\|_{\infty,1} = 1$ effectively implies that each row of \mathbf{X} has exactly one nonzero entry.

Lemma 4.4. *Let $\mathcal{X} \triangleq \{\mathbf{X} \in \{0,1\}^{d \times k} : \|\mathbf{X}\|_{\infty,1} = 1\}$. For any $d \times k$ real matrix \mathbf{L} , Algorithm 2 outputs*

$$\tilde{\mathbf{X}} = \arg \max_{\mathbf{X} \in \mathcal{X}} \text{Tr}(\mathbf{X}^\top \mathbf{L}),$$

in time $O(k \cdot d)$

Proof. By construction, each row of \mathbf{X} has exactly one nonzero entry. Let $j_i \in [k]$ denote the index of the nonzero entry in the i th row of \mathbf{X} . For any $\mathbf{X} \in \mathcal{X}$,

$$\begin{aligned} \text{Tr}(\mathbf{X}^\top \mathbf{L}) &= \sum_{j=1}^k \mathbf{x}_j^\top \mathbf{1}_j = \sum_{j=1}^k \sum_{i \in \text{supp}(\mathbf{x}_j)} 1 \cdot L_{ij} \\ &= \sum_{i=1}^d L_{ij_i} \leq \sum_{i=1}^d \max_{j \in [k]} L_{ij}. \end{aligned} \quad (26)$$

Algorithm 2 achieves equality in (26) due to the choice of j_i in line 3. Finally, the running time follows immediately from the $O(k)$ time required to determine the maximum entry of each of the d rows of \mathbf{L} . \square

C Auxiliary Lemmas

Lemma C.7. *Let a_1, \dots, a_n and b_1, \dots, b_n be $2n$ real numbers and let p and q be two numbers such that $1/p + 1/q = 1$ and $p > 1$. We have*

$$\left| \sum_{i=1}^n a_i b_i \right| \leq \left(\sum_{i=1}^n |a_i|^p \right)^{1/p} \cdot \left(\sum_{i=1}^n |b_i|^q \right)^{1/q}.$$

Lemma C.8. *For any $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{n \times k}$,*

$$|\langle \mathbf{A}, \mathbf{B} \rangle| \triangleq |\text{Tr}(\mathbf{A}^\top \mathbf{B})| \leq \|\mathbf{A}\|_F \|\mathbf{B}\|_F.$$

Proof. Treating \mathbf{A} and \mathbf{B} as vectors, the lemma follows immediately from Lemma C.7 for $p = q = 2$. \square

Lemma C.9. *For any two real matrices \mathbf{A} and \mathbf{B} of appropriate dimensions,*

$$\|\mathbf{AB}\|_F \leq \min\{\|\mathbf{A}\|_2 \|\mathbf{B}\|_F, \|\mathbf{A}\|_F \|\mathbf{B}\|_2\}.$$

Proof. Let \mathbf{b}_i denote the i th column of \mathbf{B} . Then,

$$\begin{aligned} \|\mathbf{AB}\|_F^2 &= \sum_i \|\mathbf{A}\mathbf{b}_i\|_2^2 \leq \sum_i \|\mathbf{A}\|_2^2 \|\mathbf{b}_i\|_2^2 \\ &= \|\mathbf{A}\|_2^2 \sum_i \|\mathbf{b}_i\|_2^2 = \|\mathbf{A}\|_2^2 \|\mathbf{B}\|_F^2. \end{aligned}$$

Similarly, using the previous inequality,

$$\|\mathbf{AB}\|_F^2 = \|\mathbf{B}^\top \mathbf{A}^\top\|_F^2 \leq \|\mathbf{B}^\top\|_2^2 \|\mathbf{A}^\top\|_F^2 = \|\mathbf{B}\|_2^2 \|\mathbf{A}\|_F^2.$$

The desired result follows combining the two upper bounds. \square

Lemma C.10. *For any real $m \times k$ matrix \mathbf{X} , $m \times n$ matrix \mathbf{A} , and $n \times k$ matrix \mathbf{Y} ,*

$$|\text{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y})| \leq \|\mathbf{X}\|_F \cdot \|\mathbf{A}\|_2 \cdot \|\mathbf{Y}\|_F.$$

Proof. We have

$$|\text{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y})| \leq \|\mathbf{X}\|_F \cdot \|\mathbf{A} \mathbf{Y}\|_F \leq \|\mathbf{X}\|_F \cdot \|\mathbf{A}\|_2 \cdot \|\mathbf{Y}\|_F,$$

with the first inequality following from Lemma C.8 on $|\langle \mathbf{X}, \mathbf{A} \mathbf{Y} \rangle|$ and the second from Lemma C.9. \square

Lemma C.11. *For any real $m \times n$ matrix \mathbf{A} , and pair of $m \times k$ matrix \mathbf{X} and $n \times k$ matrix \mathbf{Y} such that $\mathbf{X}^\top \mathbf{X} = \mathbf{I}_k$ and $\mathbf{Y}^\top \mathbf{Y} = \mathbf{I}_k$ with $k \leq \min\{m, n\}$, the following holds:*

$$|\text{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y})| \leq \sqrt{k} \cdot \left(\sum_{i=1}^k \sigma_i^2(\mathbf{A}) \right)^{1/2}.$$

Proof. By Lemma C.8,

$$\begin{aligned} |\langle \mathbf{X}, \mathbf{A} \mathbf{Y} \rangle| &= |\text{Tr}(\mathbf{X}^\top \mathbf{A} \mathbf{Y})| \\ &\leq \|\mathbf{X}\|_F \cdot \|\mathbf{A} \mathbf{Y}\|_F = \sqrt{k} \cdot \|\mathbf{A} \mathbf{Y}\|_F. \end{aligned}$$

where the last inequality follows from the fact that $\|\mathbf{X}\|_F^2 = \text{Tr}(\mathbf{X}^\top \mathbf{X}) = \text{Tr}(\mathbf{I}_k) = k$. Further, for any \mathbf{Y} such that $\mathbf{Y}^\top \mathbf{Y} = \mathbf{I}_k$,

$$\|\mathbf{A} \mathbf{Y}\|_F^2 \leq \max_{\substack{\hat{\mathbf{Y}} \in \mathbb{R}^{n \times k} \\ \hat{\mathbf{Y}}^\top \hat{\mathbf{Y}} = \mathbf{I}_k}} \|\mathbf{A} \hat{\mathbf{Y}}\|_F^2 = \sum_{i=1}^k \sigma_i^2(\mathbf{A}). \quad (27)$$

Combining the two inequalities, the result follows. \square

Lemma C.12. *For any real $m \times n$ matrix \mathbf{A} , and any $k \leq \min\{m, n\}$,*

$$\max_{\substack{\mathbf{Y} \in \mathbb{R}^{n \times k} \\ \mathbf{Y}^\top \mathbf{Y} = \mathbf{I}_k}} \|\mathbf{A} \mathbf{Y}\|_F = \left(\sum_{i=1}^k \sigma_i^2(\mathbf{A}) \right)^{1/2}.$$

The above equality is realized when the k columns of \mathbf{Y} coincide with the k leading right singular vectors of \mathbf{A} .

Proof. Let $\mathbf{U}\Sigma\mathbf{V}^\top$ be the singular value decomposition of \mathbf{A} , with $\Sigma_{jj} = \sigma_j$ being the j th largest singular value of \mathbf{A} , $j = 1, \dots, d$, where $d \triangleq \min\{m, n\}$. Due to the invariance of the Frobenius norm under unitary multiplication,

$$\|\mathbf{A}\mathbf{Y}\|_{\text{F}}^2 = \|\mathbf{U}\Sigma\mathbf{V}^\top\mathbf{Y}\|_{\text{F}}^2 = \|\Sigma\mathbf{V}^\top\mathbf{Y}\|_{\text{F}}^2. \quad (28)$$

Continuing from (28),

$$\begin{aligned} \|\Sigma\mathbf{V}^\top\mathbf{Y}\|_{\text{F}}^2 &= \text{Tr}(\mathbf{Y}^\top\mathbf{V}\Sigma^2\mathbf{V}^\top\mathbf{Y}) \\ &= \sum_{i=1}^k \mathbf{y}_i^\top \left(\sum_{j=1}^d \sigma_j^2 \cdot \mathbf{v}_j \mathbf{v}_j^\top \right) \mathbf{y}_i \\ &= \sum_{j=1}^d \sigma_j^2 \cdot \sum_{i=1}^k (\mathbf{v}_j^\top \mathbf{y}_i)^2. \end{aligned}$$

Let $z_j \triangleq \sum_{i=1}^k (\mathbf{v}_j^\top \mathbf{y}_i)^2$, $j = 1, \dots, d$. Note that each individual z_j satisfies

$$0 \leq z_j \triangleq \sum_{i=1}^k (\mathbf{v}_j^\top \mathbf{y}_i)^2 \leq \|\mathbf{v}_j\|^2 = 1,$$

where the last inequality follows from the fact that the columns of \mathbf{Y} are orthonormal. Further,

$$\begin{aligned} \sum_{j=1}^d z_j &= \sum_{j=1}^d \sum_{i=1}^k (\mathbf{v}_j^\top \mathbf{y}_i)^2 = \sum_{i=1}^k \sum_{j=1}^d (\mathbf{v}_j^\top \mathbf{y}_i)^2 \\ &= \sum_{i=1}^k \|\mathbf{y}_i\|^2 = k. \end{aligned}$$

Combining the above, we conclude that

$$\|\mathbf{A}\mathbf{Y}\|_{\text{F}}^2 = \sum_{j=1}^d \sigma_j^2 \cdot z_j \leq \sigma_1^2 + \dots + \sigma_k^2. \quad (29)$$

Finally, it is straightforward to verify that if $\mathbf{y}_i = \mathbf{v}_i$, $i = 1, \dots, k$, then (29) holds with equality. \square

Lemma C.13. *For any real $m \times n$ matrix \mathbf{A} , let $\sigma_i(\mathbf{A})$ be the i th largest singular value. For any $r, k \leq \min\{m, n\}$,*

$$\sum_{i=r+1}^{r+k} \sigma_i(\mathbf{A}) \leq \frac{k}{\sqrt{r+k}} \|\mathbf{A}\|_{\text{F}}.$$

Proof. By the Cauchy-Schwartz inequality,

$$\begin{aligned} \sum_{i=r+1}^{r+k} \sigma_i(\mathbf{A}) &= \sum_{i=r+1}^{r+k} |\sigma_i(\mathbf{A})| \leq \left(\sum_{i=r+1}^{r+k} \sigma_i^2(\mathbf{A}) \right)^{1/2} \|\mathbf{1}_k\|_2 \\ &= \sqrt{k} \cdot \left(\sum_{i=r+1}^{r+k} \sigma_i^2(\mathbf{A}) \right)^{1/2}. \end{aligned}$$

Note that $\sigma_{r+1}(\mathbf{A}), \dots, \sigma_{r+k}(\mathbf{A})$ are the k smallest among the $r+k$ largest singular values. Hence,

$$\begin{aligned} \sum_{i=r+1}^{r+k} \sigma_i^2(\mathbf{A}) &\leq \frac{k}{r+k} \sum_{i=1}^{r+k} \sigma_i^2(\mathbf{A}) \leq \frac{k}{r+k} \sum_{i=1}^l \sigma_i^2(\mathbf{A}) \\ &= \frac{k}{r+k} \|\mathbf{A}\|_{\text{F}}^2. \end{aligned}$$

Combining the two inequalities, the desired result follows. \square

Corollary 1. *For any real $m \times n$ matrix \mathbf{A} , the r th largest singular value $\sigma_r(\mathbf{A})$ satisfies $\sigma_r(\mathbf{A}) \leq \|\mathbf{A}\|_{\text{F}}/\sqrt{r}$.*

Proof. It follows immediately from Lemma C.13. \square

First, we define the $\|\cdot\|_{\infty,2}$ norm of a matrix as the l_2 norm of the column with the maximum l_2 norm, *i.e.*, for an $r \times k$ matrix \mathbf{C}

$$\|\mathbf{C}\|_{\infty,2} = \max_{1 \leq i \leq k} \|\mathbf{c}_i\|_2.$$

Note that

$$\begin{aligned} \|\mathbf{C}\|_{\text{F}}^2 &= \sum_{i=1}^k \|\mathbf{c}_i\|_2^2 \leq k \cdot \max_{1 \leq i \leq k} \|\mathbf{c}_i\|_2^2 \\ &= k \cdot \left(\max_{1 \leq i \leq k} \|\mathbf{c}_i\|_2 \right)^2 = k \cdot \|\mathbf{C}\|_{\infty,2}^2. \end{aligned} \quad (30)$$