

Learning to Partition using Score Based Compatibilities

Arun Rajkumar
Conduent Labs - India
arun.rajkumar@conduent.com

Koyel Mukherjee
IBM Research - India
kmukherj@in.ibm.com

Theja Tulabandhula
Univ. of Illinois Chicago
tt@theja.org

ABSTRACT

We study the problem of learning to partition users into groups, where one must learn the compatibilities between the users to achieve optimal groupings. We define four natural objectives that optimize for average and worst case compatibilities and propose new algorithms for adaptively learning optimal groupings. When we do not impose any structure on the compatibilities, we show that the group formation objectives considered are *NP* hard to solve and we either give approximation guarantees or prove inapproximability results. We then introduce an elegant structure, namely that of *intrinsic scores*, that makes many of these problems polynomial time solvable. We explicitly characterize the optimal groupings under this structure and show that the optimal solutions are related to *homophilous* and *heterophilous* partitions, well-studied in the psychology literature. For one of the four objectives, we show *NP* hardness under the score structure and give a $\frac{1}{2}$ approximation algorithm for which no constant approximation was known thus far. Finally, under the score structure, we propose an online low sample complexity PAC algorithm for learning the optimal partition. We demonstrate the efficacy of the proposed algorithm on synthetic and real world datasets.

1. INTRODUCTION

The problem of learning to partition users (or objects) into groups has numerous applications in ridesharing, health care groups, project groups etc. Effective grouping of users is critical in these applications as it determines how well the group can work together [27]. For instance, the success of ridesharing adoption depends on how the system assigns users to groups taking into account their source destination requirements as well as their personalities. A user who does not like making conversation would rate the ride-sharing experience low if they are in the same group as users who talk a lot. Here, the ridesharing company might be interested in maximizing the average *happiness* of the rideshare groups. In another instance, effectively allocating employees to teams for a critical project might influence the project's success [21]. Here the organization might want to be defensive and maximize the compatibility of the least compatible pair of employees in any team.

Appears in: *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2017)*, S. Das, E. Durfee, K. Larson, M. Winikoff (eds.), May 8–12, 2017, São Paulo, Brazil.
Copyright © 2017, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

That effective grouping leads to improved outcomes has been extensively studied in the psychology literature [18]. In particular, the effects of *homophily* (grouping similar individuals together) [23, 2] and *heterophily* (grouping dissimilar individuals together) [25] on desired outcomes have been documented in applications such as team formation [19] and study groups [10, 13] among others. While the social phenomenon of homophily and heterophily have been well studied [20] [1], it is not clear (more so from a formal standpoint) as to when one should prefer homophilous to heterophilous grouping in general.

In this work, we formalize the group formation problem by considering four concrete objectives that involve pairwise compatibility of users. We refer to the pairwise compatibility as the *happiness* of the pair of users. Under this happiness index, we study the following objectives: maximize (i) the average of average *happiness* across groups (AoA), (ii) minimum compatibility of any pair of users across groups (MoM), (iii) average of the minimum compatible pairs across groups (AoM) and (iv) minimum of the average happiness across all groups (MoA) (see Table 2 for formal definitions). The objectives cover most of the scenarios one may be interested in maximizing. While we do consider general pairwise compatibilities, we introduce and focus mainly on an elegant structure that we impose on them. Under this structure, each individual has an associated *intrinsic score* and the pairwise compatibility between two individuals depend on their corresponding scores. We will see that under this structure, the optimal solutions to these objectives naturally translate to homophilous and heterophilous groupings. The motivation of studying the score-based model, where the compatibility (or, preference towards each other) of a pair of users is determined by their relative scores, arises from the well-known Bradley-Terry-Luce (BTL) (Bradley and Terry, 1952 [5]; Luce, 1959 [22]) in literature. The BTL model is often applied for pairwise comparisons, in order to determine the relative preferences. This led us to study the simpler case of score-based preferences, for which optimal polynomial algorithms can be provided in most cases, while the general case is either inapproximable or hard to approximate beyond a certain factor.

We list below the major contributions of this paper.

Our Contributions

- We show that all the objectives considered above are NP-hard (and hard to approximate beyond a factor) when we assume no structure on the pairwise compatibilities. We give polynomial algorithms under certain assumptions.

Table 1: Summary of main results of the paper. All shaded boxes correspond to NP-Hard problems. The top row (in blue) lists the hardness results, and the bottom row highlights the polynomial time solvability (unshaded) and approximation guarantee (gray).

	AoA	MoM	AoM	MoA
General	$1 - \frac{1}{\binom{k}{2}^m}$	INAPPROX	$(1 - \frac{1}{m})$	$1 - \frac{1}{\binom{k}{2}}$
Score-based	Homophily	Heterophily	Homophily	$\frac{1}{2}$

- When the compatibilities are dictated by intrinsic scores of each individual, we show that three of the four objectives become solvable in polynomial time.
- Under the intrinsic scores assumption, we explicitly characterize the structure of the optimal solution and show that they are related to the notion of homophily and heterophily.
- We show that the MoA objective is NP-hard to solve even under the intrinsic scores assumption.
- We propose a greedy algorithm for the MoA objective and prove a $\frac{1}{2}$ approximation guarantee for the same. This is a significant improvement on known algorithms, since no constant factor result was known thus far.
- Under the intrinsic score structure, we propose a low sample complexity PAC learning algorithm to learn the scores (and hence the optimal groupings) provably.

Table 1 summarizes some of the contributions of this paper.

2. RELATED WORK

Graph partitioning and clustering problems are well researched with several variants such as the disjoint clique problem, k-way partition problem, partition into triangles and many others [12]. Balanced or capacitated versions of such partitioning or clustering problems are related to our work. For instance, the AoA objective is the same as the k-equipartition problem [11] and is known to be NP-hard [14]. The SUM-MAX partitioning problem [28] looks at minimizing the average of maximum weighted edge between pairs of groups and is close to, but not the same as the AoM objective. Min-max objectives similar to MoM and MoA objectives have been recently considered by Svitkina and Tardos [26] and Bansal et al. [3], where the problem is to create a k-way equipartition to minimize the maximum weight of edges leaving a single group. This is different from the MoM objective where we want to maximize the minimum of minimum weights across edges within each group.

Many seemingly related problems such as the clique partition problem [16], the capacitated max-k-cut and min-k-cut problems [15] do not come with size restrictions or have objectives that do not capture homophily or heterophily. In the clustering domain, the work closest to ours is that of Zhu, Wang and Li [29], where heuristic algorithms for clustering with size constraints on the clusters are proposed, although

(AoA)	$\max_{\Pi(\mathbf{S}_1, \dots, \mathbf{S}_m)} \frac{1}{m} \sum_{i=1}^m H(\mathbf{S}_i \mathbf{W})$
(MoM)	$\max_{\Pi(\mathbf{S}_1, \dots, \mathbf{S}_m)} \min_i \left(\min_{j, k \in \mathbf{S}_i} W_{jk} \right)$
(AoM)	$\max_{\Pi(\mathbf{S}_1, \dots, \mathbf{S}_m)} \frac{1}{m} \sum_{i=1}^m \left(\min_{j, k \in \mathbf{S}_i} W_{jk} \right)$
(MoA)	$\max_{\Pi(\mathbf{S}_1, \dots, \mathbf{S}_m)} \min_i H(\mathbf{S}_i \mathbf{W})$

Table 2: Objectives considered.

for a majority of clustering applications such size constraints or information on the number of clusters is not pre-defined.

There is one class of graph partitioning problems that come with size restrictions, namely, the Graph Tree Partitioning Problem [7], where the objective is to partition a graph into equal size subsets, such that the weight of the spanning tree on each subset is either as high as possible (max-min) or as low as possible (min-max), or, the sum of the weights of the spanning trees are either as high (max-sum), or, as low (min-sum) as possible. Though these objectives are closely related to ours, they are not exactly the same problems.

Some recent research in application areas such as team formation and ride-sharing study similar questions as us. Singla et al. [24] present online learning algorithms with PAC bounds in the context of learning the expertise of workers for team formation. The emphasis is on speeding up learning by exploiting similarities between workers and between tasks. In contrast to our work, their objective is to select an optimal subset of the workers after learning their expertise and there is no notion of forming groups. Bistaffa, Farinelli and Ramchurn [4] study an offline optimization problem of grouping users to minimize travel costs. This is formulated as a coalition formation problem, restricted by a social network based compatibility graph that is assumed given. The groups are formed based on heuristics without any guarantees on optimality. Brindley, Blaschke and Walti [6] look at what factors impact the formation of effective learning groups through an empirical study.

Our formulations and learning algorithms can also be applied in recurring sharing economy settings (e.g., AirBnB) as well as healthcare. In the latter setting, it has been observed that assigning patients of similar disease characteristics to groups often helps in effective treatment [8].

3. PRELIMINARIES

Let $[n] = \{1, 2, \dots, n\}$ be the set of items to be partitioned into groups. Let each group be of size k and let $m = \frac{n}{k}$ denote the number of groups¹. A k -partition $\Pi(\mathbf{S}_1 \dots \mathbf{S}_m)$ of $[n]$ is denoted by a set of subsets $\{\mathbf{S}_1, \mathbf{S}_2 \dots \mathbf{S}_m\}$ where each $\mathbf{S}_i \subseteq [n]$, $|\mathbf{S}_i| = k$ with $\mathbf{S}_i \cap \mathbf{S}_j = \emptyset \forall i \neq j$ and $\bigcup_i \mathbf{S}_i = [n]$. We will capture the relation between users/objects using a symmetric *pairwise compatibility matrix* $\mathbf{W} \in \mathbf{R}_+^{n \times n}$ where $W_{ij} (= W_{ji})$ denotes the compatibility of users/objects i and j . Given a subset $\mathbf{S} \subseteq [n]$, we define the *happiness* of the subset with respect to \mathbf{W} as $H(\mathbf{S} | \mathbf{W}) = \frac{1}{|\mathbf{S}|^2} \sum_{i, j \in \mathbf{S}} W_{ij}$.

¹Assume k divides n or add *dummy* items if not.

Problem Definition: Given a pairwise compatibility matrix $\mathbf{W} \in \mathbb{R}_+^{n \times n}$, partition the n items into m groups in the best possible manner that maximizes each of the four objectives defined in Table 2.

Towards this, we will first consider the case where we don't impose any conditions on the pairwise compatibilities. In the subsequent section, we consider the same problem by imposing a score structure.

4. GENERAL COMPATIBILITY MATRICES

We start by describing some results on the hardness of approximation of the four objectives in the general case.

THEOREM 1. (Approximability) For $k \geq 3$, unless $P = NP$, the following are lower bounds on polynomial time approximability: (a) MoM: inapproximable, (b) AoM: $(1 - \frac{1}{m})$, (c) AoA: $1 - \frac{1}{\binom{k}{2}^m}$ and (d) MoA: $1 - \frac{1}{\binom{k}{2}}$.

PROOF. MoM: inapproximable. Consider an instance $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ of PARTITIONINTOTRIANGLES [14], where $|\mathcal{V}| = 3q$ for some $q \in \mathbb{Z}^+$. The decision question is whether \mathcal{V} can be partitioned into q disjoint sets of size 3 each, V_1, \dots, V_q , such that each $V_i, i \in [q]$ forms a triangle in \mathcal{G} . Create an instance of MoM, a weighted graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}', \mathcal{W}')$ from \mathcal{G} , where for every vertex $v \in \mathcal{V}$, we create a vertex $v' \in \mathcal{V}'$, and for every edge $e = (u, v) \in \mathcal{E}$, we create an edge $e' = (v', u') \in \mathcal{E}'$ between the corresponding vertices v' and u' in \mathcal{V}' , and set its weight $w_{e'} = M$, where M is a large number. Set $k = 3$ (group size) and $m = q$ (number of groups). For any pair of vertices $(p, q) \in \mathcal{V}$, such that no edge exists between p and q in \mathcal{E} , we create an edge $e'' = (p, q) \in \mathcal{E}'$ of weight $w_{e''} = \epsilon$, where ϵ is a small number and add it to \mathcal{E}' . If there exists a partition of \mathcal{V} into q disjoint triangles, then there exists a solution to the MoM problem with objective function $= M$. The q disjoint triangles correspond to m groups in \mathcal{G}' , where in each group, every edge has weight M , since the corresponding edge exists in \mathcal{G} . Similarly, if there does not exist any partition of \mathcal{V} into q disjoint triangles, then every solution to MoM on \mathcal{G}' has value $= \epsilon$. This is because any solution to the group formation problem with $k = 3$ and $m = q$ on \mathcal{G}' , would result in at least one group where at least one edge has a weight ϵ , since the corresponding edge does not exist in \mathcal{G} .

Therefore, if there exists a polynomial time approximation algorithm with an approximation ratio $> \frac{\epsilon}{M}$, one would be able to distinguish between the YES and NO instances of PARTITIONINTOTRIANGLES, by creating a corresponding instance of MoM and applying the algorithm. A YES instance would result in MoM $> \epsilon$, and similarly, a NO instance would result in MoM $\leq \epsilon$. Hence, unless $P = NP$, there can be no polynomial time approximation algorithm with an approximation ratio $> \frac{\epsilon}{M}$. For $\epsilon \rightarrow 0$, and $M \rightarrow \infty$, we can make the ratio arbitrarily bad, $\rightarrow 0$. Hence, it is NP-hard to approximate the MoM problem in the general case.

The approximability proofs of **MoA**, **AoA** and **AoM** use the same reduction as above. Specifically, create a graph $\mathcal{G}' = (\mathcal{V}', \mathcal{E}', \mathcal{W}')$, with $k = 3$ and $m = q$, from an instance of PARTITIONINTOTRIANGLES, $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $|\mathcal{V}| = 3q$. For a YES instance of the PARTITIONINTOTRIANGLES, the disjoint m groups in \mathcal{G}' corresponding to the disjoint q triangles in \mathcal{G} , will give a solution with MoA=AoA=AoM= M . On the

other hand, for a NO instance, any partition of \mathcal{G}' into m disjoint groups would result in at least one group, with at least one edge of weight ϵ . Setting $\epsilon = 0$, therefore, MoA $\leq \frac{2M}{3}$, AoA $\leq M(1 - \frac{1}{3m})$, and AoM $\leq M(1 - \frac{1}{m})$.

Therefore, if there exists a polynomial time approximation algorithm for any of the above three objectives with better approximation factors (MoA: $\frac{2}{3}$, AoA: $(1 - \frac{1}{3m})$, and AoM: $(1 - \frac{1}{m})$), one can distinguish between the YES and NO instances of PARTITIONINTOTRIANGLES. The approximation guarantees can be extended to general $k > 3$, by similar reduction, replacing triangles by k -cliques.

□

4.1 Polynomial Solvability for $k = 2$

While the general case is hard, for the case where $k = 2$, all the objectives become polynomial time solvable.

THEOREM 1. (Polynomial solvability for $k = 2$) When $k = 2$, all four objectives are polynomial time solvable.

For proving Theorem 1, we first prove the following claim.

CLAIM 2. When $k = 2$, the optimal solution for MOA is the same as that of MOM, and the optimal solution for AoA is the same as the optimal solution for AoM.

PROOF. We first prove that the optimal solution for MoA is the same as that for MoM when $k = 2$. Since group sizes are $k = 2$, every group has only one edge, that occurring between the pair of vertices in the group. Therefore, the average of the weight of edges in any group is determined simply by the weight of the single edge in the group. Hence, the minimum of the averages, namely MoA is the same as the minimum weight edge in any group, that is, MoM. We next argue that the optimal solution for AoA is the same as that for AoM when $k = 2$. Since every group has only a single edge each, the solution maximizing the average of the average weight of every group, namely, AoA, is the same as maximizing the average of the weight of the edge in every group, which corresponds to AoM for $k = 2$. □

Now, we discuss how Edmond's maximum weighed matching for general graphs solves AoA (and AoM) optimally in polynomial time for $k = 2$. We find a maximum weight matching in \mathcal{G} using Edmond's algorithm [9]. This returns a partition of \mathcal{V} into $\frac{n}{2}$ groups, maximizing the total weight of edges used in the matching. Suppose there is an optimal solution for AoA (also, AoM) for $k = 2$ in \mathcal{G} , that has a higher objective function value. That means that there exists a partition of \mathcal{V} into $\frac{n}{k}$ disjoint subsets, with each subset of size 2, hence including only one edge, such that the total weight of edges used in the subsets is higher than that returned by Edmond's algorithm. However, that contradicts the optimality of Edmond's algorithm, since we can use the solution to AoA (also, AoM) to find a higher weight matching. Similarly, suppose there is a higher weight matching than the objective function value of the optimal solution for AoA (also, AoM). Then, one can use the matching returned by Edmond's algorithm to find a higher value, feasible solution for AoA (also, AoM), thereby contradicting the optimality of the optimal value for AoA (also, AoM).

For MoA and MoM, we first prove the following property.

LEMMA 2. Given a graph \mathcal{G} , and an optimal solution value OPT for MoM (also, MoA), if we delete all edges of weight

$< OPT$, and make the resultant graph \mathcal{G}' unweighted, then there exists a perfect matching in \mathcal{G}' , corresponding to an optimal solution in \mathcal{G} .

PROOF. Consider the groups in an optimal solution for MoM (also, MoA) in \mathcal{G} . Since the optimal solution has value OPT , there must exist a partition of the vertices in to groups of 2, such that the edge in each group has a weight $\geq OPT$. However, this corresponds to a matching in \mathcal{G} , where every edge in the matching has a weight $\geq OPT$. Now, delete every edge of weight $< OPT$. The matching remains unperturbed, and corresponds to a perfect matching in \mathcal{G}' . \square

Now, we give the algorithm for optimally solving MoM and MoA for $k = 2$. The algorithm for maximizing MoM (also, MoA) would involve ordering the distinct weights in \mathcal{W} in non-decreasing order. For the next weight w in the list, delete all edges of weight $w_e < w$. Create an unweighted graph \mathcal{G}' containing only the remaining edges in \mathcal{G} (without considering their weights). Now, use Edmond's algorithm to find a maximum cardinality matching in \mathcal{G}' . The lowest weight w , after deleting which, \mathcal{G}' does not have a perfect matching, is the optimal value for MoM (also, MoA). The optimality of this algorithm follows from Lemma 2 and the optimality of Edmond's maximum cardinality matching for general graphs, that runs in polynomial time. The polynomial time solvability also follows, hence.

We next study a linear time solvable special case.

4.2 Transitive Compatibility: Optimal Linear Algorithm for any k

We next prove that under a transitivity like assumption on the compatibility matrix, there exists a linear time optimal algorithm. The assumption follows from the intuition that if user i is compatible with j , j is compatible with k , then i is compatible with k . Formally, the transitive property assumed is: $\forall i, j, k, W_{ij} \geq \min(W_{ik}, W_{kj})$. The following theorem follows from the fact that graphs obeying transitive compatibility would have a particular structure: a collection of disjoint cliques, and, a linear traversal of the graph would return the optimal solution.

THEOREM 3. *Under the transitive compatibility property in \mathcal{G} , there exists a linear time optimal algorithm for MoM.*

Proof Sketch:

We first argue that the graph \mathcal{G} on which the transitive compatibility property holds has a certain structure. Specifically, it consists of disjoint connected components, where each connected component is a clique. To see this, consider a pair of vertices in \mathcal{G} , u and v , between whom the happiness or compatibility is 0. In other words, no edge exists between them. We argue that in order to maintain the transitive property, any vertex p that u is adjacent to (i.e., $W_{u,p} > 0$), must necessarily have a 0 compatibility with v . This is because if there exists an edge between p and v , then with both $W_{p,v} > 0$ and $W_{u,p} > 0$, the transitive property is violated by $0 = W_{u,v} < \min(W_{u,p}, W_{p,v})$. Therefore, if $W_{u,v} = 0$, then $W_{p,v} = 0$ for all p , such that $W_{u,p} > 0$. Similarly, any vertex p' that p is adjacent to cannot have any edge to v . Continuing in this manner, it can be seen that any vertex that u is connected to cannot have any edge to v . Therefore, u and v must be in disjoint connected components. Similarly, any pair of vertices that u is adjacent to, say, v_1 and v_2 , must

be adjacent to each other. Otherwise, it can be seen that the transitive compatibility property is violated.

Another property that transitive compatibility induces is that, in every clique, there can be at most one edge of higher weight, and all other edges must be of identical weight. To see this, note that if there are two edges of higher weight compared to the weight of all other edges in a clique, then there would be at least one triangle, where one edge is lower in weight than the other two, violating the transitive property.

Any optimal solution would consider each clique separately (in other words, only the participants belonging to the same clique would be matched to one another), since otherwise the partitions would include edges of weight = 0. Moreover, in every grouping of the clique vertices, there can be at most one group with at most one edge of higher weight, and all other edges in all the groups would have identical lower weight, say w . Hence, replacing the higher weight edge by an edge of weight w would not change the MoM objective value. Hence, any partitioning of the vertices of a clique in to groups of size k^2 would be optimal. The algorithm is linear in the number of vertices, since one has to traverse at most all the vertices in all the cliques in order to get the partitions.

5. SCORE BASED COMPATIBILITY MATRICES

In this section we consider a simple yet useful structure on the pairwise compatibility matrix under which three out of the four objectives introduced in Section 3 become poly-time solvable. Specifically, we consider the case of *score based compatibility* matrices where every item i has an associated score $s_i \in \mathbf{R}_+$ and the pairwise compatibility of items i and j is given by the product of their individual scores $s_i s_j$. This is a natural assumption in several practical scenarios. For instance in study groups, the score could refer to the influence a student has on her peers which may depend on the GPA of the student. In ride sharing applications the score may indicate how introverted/extroverted a person is.

We begin by defining certain natural partitions induced by score vectors.

DEFINITION 4. (Homophilous Partition) *Let $\mathbf{s} \in \mathbf{R}^n$ and let $\sigma = \text{argsort}(\mathbf{s})^3$. A k -partition $\Pi(\mathbf{S}_1, \dots, \mathbf{S}_m)$ corresponding to \mathbf{s} is called homophilous w.r.t \mathbf{s} if $\forall i \in [m], \mathbf{S}_i = \{\sigma((i-1)k+1), \dots, \sigma((i-1)k+k)\}$*

DEFINITION 5. (Heterophilous Partition) *Let $\mathbf{s} \in \mathbf{R}^n$ and let $\sigma = \text{argsort}(\mathbf{s})$. A k -partition $\Pi(\mathbf{S}_1, \dots, \mathbf{S}_m)$ corresponding to \mathbf{s} is called heterophilous if $\forall i \in [m], \mathbf{S}_i = \{\sigma((i-1)(k-1)+1), \dots, \sigma((i-1)(k-1)+(k-1)), \sigma(n+1-i)\}$*

As an example, let $\sigma = (1\ 2\ 3\ 4\ 5\ 6)$ and $k = 2$. The homophilous 2-partition corresponding to σ would be $\{\mathbf{S}_1 = (1\ 2), \mathbf{S}_2 = (3\ 4), \mathbf{S}_3 = (5\ 6)\}$ whereas the heterophilous 2-partition would be $\{\mathbf{S}_1 = (1\ 6), \mathbf{S}_2 = (2\ 5), \mathbf{S}_3 = (3, 4)\}$. Our main results of this section explicitly characterize the optimal solutions for the objectives considered.

²Without loss of generality, for any clique of size n' , we add dummy vertices with edges of weight w incident on them, to make its cardinality, a multiple of k .

³ $\text{argsort}(\mathbf{s})$ is the permutation obtained by sorting the values of score vector \mathbf{s} in non-increasing order. Specifically, for $\sigma = \text{argsort}(\mathbf{s})$, for any $i, j, s_i > s_j \implies \sigma(i) < \sigma(j)$.

THEOREM 6. (Homophilous Partition is Optimal for AoA and AoM) Let $\mathbf{W} \in \mathbb{R}^{n \times n}$ be a score based compatibility matrix parametrized by the score vector $\mathbf{s} \in \mathbb{R}^n$. The optimal solution to the average of averages (AoA) and the average of minimums (AoM) objectives w.r.t \mathbf{W} is given by the homophilous partition of \mathbf{s} .

PROOF. Note that for any group \mathbf{S} , if the compatibility matrix is score based, then the sum of weights of all pairs in the group is given by $\sum_{i,j \in \mathbf{S}} W_{ij} = \left(\sum_{i \in \mathbf{S}} s_i \right)^2$

AoA Objective: Assume wlog the entries of \mathbf{s} are sorted in descending order i.e., $s_i > s_j$ for all $i < j$. Let Π denote the homophilous k -partition corresponding to \mathbf{s} . For the sake of contradiction let $\bar{\Pi} \neq \Pi$ be the optimal k -partition. Let i be the minimum index such that both i and $i+1$ are in the same groups in Π whereas they are in different groups in $\bar{\Pi}$. Denote these groups by g_1 and g_2 . We will show that by swapping specific elements from g_1 and g_2 , one can obtain a partition which is at least as good as $\bar{\Pi}$. Let $i \in g_1$ and $i+1 \in g_2$. Notice that there must be at least one element $x \in g_1$ such that $s_{i+1} > s_x$ (otherwise it contradicts the minimality of i). We will consider two cases depending on whether $\sum_{j \in g_1} s_j > \sum_{j \in g_2} s_j$ or otherwise. Case 1: $\sum_{j \in g_1} s_j > \sum_{j \in g_2} s_j$. Let $v_1 = \sum_{j \in g_1: j \neq i, x} s_j$, $v_2 = \sum_{j \in g_2: j \neq i+1} s_j$. Then, using convexity of $H(\mathbf{S}_i | \mathbf{W})$, we have

$$(s_i + s_{i+1} + v_1)^2 + (s_x + v_2)^2 \geq (s_i + s_x + v_1)^2 + (s_{i+1} + v_2)^2$$

Case 2: $\sum_{j \in g_1} s_j \leq \sum_{j \in g_2} s_j$ In this case, there must be an element $x \in g_2$ such that $s_x < s_i$. One can follow a similar proof as the previous case by swapping $i \in g_1$ and $x \in g_2$. In both the cases, we obtain a partition whose sum (average) of weights over the groups of the partition is at least as good as $\bar{\Pi}$. One can repeat the procedure with the new partition obtained until one reaches Π . But this contradicts the fact that $\bar{\Pi} \neq \Pi$ is optimal.

AoM Objective: Assume wlog that the score vector \mathbf{s} is such that $s_1 \leq s_2 \leq \dots \leq s_n$. We will show the result using induction on the number of groups m . Consider the base case where $m = 2$ i.e. $n = 2k$. In this case, we need to show:

$$s_1 s_2 + s_{k+1} s_{k+2} \geq s_{i_1} s_{j_1} + s_{i_2} s_{j_2}$$

where (i_1, j_1) and (i_2, j_2) correspond to the minimum compatible pairs in the two groups corresponding to some non-homophilous partition of \mathbf{s} . If any of these pairs is same as $(1, 2)$ then the result is obvious. Assume not. Then both 1 and 2 will contribute to the minimum compatible pairs. Thus it is enough to show that both of the below cases hold

$$s_1 s_2 + s_{k+1} s_{k+2} \geq s_1 s_{k+2} + s_2 s_3 \text{ and}$$

$$s_1 s_2 + s_{k+1} s_{k+2} \geq s_1 s_3 + s_2 s_{k+2}$$

as all other cases give rise to smaller objective values. But

$$\begin{aligned} s_{k+2}(s_3 - s_2) &\geq s_1(s_3 - s_2) \\ \implies s_1 s_2 + s_{k+1} s_{k+2} &\geq s_1 s_{k+2} + s_2 s_3 \end{aligned}$$

Similarly one can show the result for the other case as well. This proves the base case. Now for a general k , assume that the induction hypothesis is true for $m = k - 1$. For $n = mk$, apply the induction hypothesis to the bottom $n - k$ items i.e, for the set $\{s_{k+1}, \dots, s_n\}$. We need to show that with the newly added items $\{s_1, s_2, \dots, s_k\}$ the hypothesis is still

satisfied. Assume not. Then there must exist index pairs $(i_1, j_1), \dots, (i_m, j_m)$ such that

$$\begin{aligned} s_1 s_2 + s_{k+1} s_{k+2} + \dots + s_{(m-1)k+1} s_{(m-1)k+2} < \\ s_{i_1} s_{j_1} + s_{i_2} s_{j_2} + \dots + s_{i_m} s_{j_m} \end{aligned}$$

If items 1 and 2 are in the same group, we arrive at a contradiction. Assume they are in different groups. Let $(1, j_1)$ and $(2, j_2)$ be the corresponding minimum pairs. By swapping 2 with j_1 , we can only increase the objective. We can iteratively swap items without decreasing the objective such that the first k items are in the first group. But this contradicts the induction hypothesis for the last $n - k$ items. \square

THEOREM 7. (Heterophilous Partition is Optimal for MoM) Let $\mathbf{W} \in \mathbb{R}^{n \times n}$ be a score based compatibility matrix parametrized by the score vector $\mathbf{s} \in \mathbb{R}^n$. The optimal solution to the minimum of minimums (MoM) objective w.r.t \mathbf{W} is given by the heterophilous partition of \mathbf{s} .

PROOF. We prove the base case for $n = 4$ and $k = 2$. Let the scores be given by $s_1 > s_2 > s_3 > s_4$. The three possible ways of partitioning this are given by $\{(1, 2)(3, 4)\}, \{(1, 3)(2, 4)\}, \{(1, 4)(2, 3)\}$. Note that we have $s_1 s_4 > s_3 s_4$ and $s_2 s_3 > s_3 s_4$. Thus,

$$\min(s_1 s_4, s_2 s_3) > s_3 s_4 = \min(s_1 s_2, s_3 s_4)$$

Similarly, we have $s_1 s_4 > s_2 s_4$ and $s_2 s_3 > s_2 s_4$. Thus,

$$\min(s_1 s_4, s_2 s_3) > s_2 s_4 = \min(s_1 s_3, s_2 s_4)$$

As the induction hypothesis, assume that the claim is true for some n and k . We will show that it is true for $n + k$. Let the new items added be x_1, x_2, \dots, x_k assume wlog that $x_1 > x_2 > x_{k-1} > s_1 > s_2 > \dots > s_n > x_k$. We know from the induction hypothesis that the heterophilous partition corresponding to \mathbf{s} is has the highest MoM objective value. For the sake of contradiction, assume that by adding the new k items, the heterophilous partition corresponding to the vector $[x_1 \ x_2 \ \dots \ x_{k-1} \ s_1 \ s_2 \ \dots \ s_n \ x_k]$ is not optimal.

Case 1: As x_1, \dots, x_{k-1} are larger than $s_i \ \forall i$ and x_k , the only way any of them could be a part of the minimum pair is when all the items with scores x_1, \dots, x_{k-1} are in the same group. If x_k is also in this group, we arrive at a contradiction. If item with score x_k is not in this group, then we swap the item with some item with score s_p to arrive at a contradiction.

Case 2: The other case to consider is when none of the items with scores $\{x_1, \dots, x_{k-1}\}$ contribute to the minimum compatible pair in their respective groups. In this case, we start with the group g that contains at least one item with scores of $\{x_1, \dots, x_{k-1}\}$ and has the highest minimum compatible weight. We iteratively swap the remaining items with scores x_i for some i which are not in g , with items from g with score s_j for some j . This can be done without decreasing the MoM objective value. We then continue swapping iteratively to get to the partition where the group g consists of items with scores $\{x_1, \dots, x_{k-1}\}$ along with some other item with score s_p for some p . Now, we can use the argument from case 1 to arrive at the required contradiction. \square

THEOREM 8. (Hardness of MoA) Computing the optimal partition for the minimum of averages (MoA) objective for a score based compatibility matrix is NP-hard.

PROOF. Consider an instance \mathcal{I} of 3-PARTITION, with $n = 3m$ items, each associated with a value; s_i is the value for the item i . The total sum of the values is $\sum_{i \in \mathcal{I}} s_i = mB$,

and the size of each item is $\frac{B}{4} < s_i < \frac{B}{2}$. The decision problem is whether there exists a partition of the items into m partitions, such that the sum of the values of items in each partition is exactly B . This is a strongly NP-hard problem.

Now, construct an instance \mathcal{I}' of the MoA problem, with $k = 3$, where we create an item in \mathcal{I}' corresponding to every item in \mathcal{I} , and the score associated with the item in \mathcal{I}' is set to the value of the corresponding item in \mathcal{I} . In the MoA problem our goal is to partition the items in \mathcal{I}' into m groups, each containing exactly $k = 3$ items, such that the total score of each group is as high as possible. In fact, the lowest total score determines MoA. The decision question we ask here is as follows: does there exist a partitioning of the n items into m groups, each group containing exactly $k = 3$ items, such that the total score of each group is at least B ?

If there exists a partition of the items into m groups, such that every group contains exactly 3 items, and the total score of every group is $\geq B$, then that corresponds to a YES instance of 3-PARTITION. Clearly, every group has to sum up to exactly B , since the total sum is mB . Alternately, if \mathcal{I} corresponds to a YES instance of 3-PARTITION for a given value B , then note that the corresponding m partitions would have exactly 3 items because of the choice of the range of the values of the items (and hence, the range of the scores), and each partition would sum up to B . This would give a feasible solution for the MoA in \mathcal{I}' , with each group's score summing up to $\geq B$. This completes the reduction. \square

While MoA is NP-Hard in general, we give a simple algorithm GREEDY for the MoA objective given $\mathbf{W} \in \mathbb{R}^{n \times n}$, a score based compatibility matrix: Sort the objects/items by their scores in a non-increasing order. Take the next unassigned item from the list and assign it to the partition (or group) with the lowest total score thus far, as long as the partition is not full (i.e., it has $< k$ vertices). Break ties arbitrarily. The best known approximation factor for this algorithm is $\max(\frac{2}{k}, \frac{1}{m})$ ([17]). We prove a constant factor approximation below, a significant improvement.

THEOREM 9. (Greedy Algorithm is $\frac{1}{2}$ approx for MoA) Algorithm GREEDY produces a k -partition that is a $\frac{1}{2}$ approximation for the MoA objective for score based compatibilities.

PROOF. For the first m iterations, each group will receive one item each from the top of the sorted list. Define a Reduced Set \mathcal{R} as a set of groups in the final solution obtained by the greedy algorithm, such that: (a) for any $(p, q) \in \mathcal{R}$, there is at least one iteration t of the greedy algorithm after the first m iterations when p is favored over q for assigning the next available item, while q is not full, and there is also at least one iteration $t' > m$ when q receives an item while p is not full, (b) \mathcal{R} has the maximum cardinality among all such sets of groups, (c) \mathcal{R} includes the group $j \in [m]$ that receives the m^{th} item in the sorted list as its first item.

Let $v_{j,i}$ be the i^{th} item added to the j^{th} group with size $s_{j,i}$, and the earliest iteration by which all groups in \mathcal{R} get full be t_R . Then it follows from the definition of \mathcal{R} :

OBSERVATION 3. Any group $j \notin \mathcal{R}$ would receive all the items from from the second one to the k^{th} one, that is, items $\{v_{j,2}, \dots, v_{j,k}\}$ in iterations $t > t_R$.

This follows from the definition of \mathcal{R} . Suppose a group $j' \notin \mathcal{R}$ received the item $v_{j',2}$ in an iteration $t' < t_R$. In that case, there exists at least one iteration, specifically, t' , when

j' is favored over each of the groups in \mathcal{R} , and at least one iteration $\leq t_R$, when each of the groups in \mathcal{R} are favored over j' , since by t_R all groups in \mathcal{R} get full. This implies, that j' should have been included in \mathcal{R} , and \mathcal{R} is not a maximal set.

CLAIM 10. Let the items in the sorted list be $\{v_1, v_2, \dots, v_n\}$, where $s_1 \geq s_2 \geq \dots \geq s_n$, (s_i is the score of item v_i , the i^{th} item in the sorted list). Let $m' = m - |\mathcal{R}|$. The items $v_{m'+1}, v_{m'+2}, \dots, v_{m'+|\mathcal{R}|k}$ get assigned to groups in \mathcal{R} .

PROOF. Suppose an item in the sequence $v_{m'+r}$, $r \in \{1, \dots, |\mathcal{R}|k\}$ gets assigned to a group $j \notin \mathcal{R}$.

Case 1: $v_{m'+r}$, $r \geq 1$ is the first item assigned to j . Now, there is at least one group $j' \in \mathcal{R}$, that receives an item earlier in the list than $v_{m'+r}$ as its first item as the greedy algorithm assigns one item each to each of the groups, before assigning the second item to any group, and there are only m' groups outside \mathcal{R} . Thus, once j' receives the second item, the total score of j' will be greater than that of j , and hence, j will receive at least one item before j' is chosen again. But from Observation 3, groups $\notin \mathcal{R}$ do not receive their second item till all groups in \mathcal{R} are full.

Case 2: $v_{m'+r}$ is the second item assigned to $j \notin \mathcal{R}$. Since from Observation 3, j receives its second item only after t_R , and all groups receive at least one item before any group receives its second item, this is possible only if $m' + r > |\mathcal{R}|k + m'$. However, $m' + r \leq |\mathcal{R}|k + m'$ by assumption. Hence this is not possible either. \square

From Claim 10, it follows that the m' groups $\notin \mathcal{R}$, receive the m' largest score items as their first items, specifically, items $v_1, \dots, v_{m'}$. Let us call these first m' items as large items, since the groups receiving them do not get their second item till all groups in \mathcal{R} are full. Moreover, it also follows that any subset of $|\mathcal{R}|k$ items from $\mathcal{V} \setminus \{v_1, \dots, v_{m'}\}$ (that is, from the set of all items excluding the large items), would have a total score at most the total score of the items used to fill \mathcal{R} . This can be seen from the fact that \mathcal{R} gets the highest score $|\mathcal{R}|k$ items, excluding the large items.

CLAIM 11. Denote the sum of scores of a set of items \mathcal{S} as $Sum(\mathcal{S})$. $\max_{(\mathcal{S} \subset \mathcal{V} \setminus \{v_1, \dots, v_{m'}\}) \cap (|\mathcal{S}| = |\mathcal{R}|k)} Sum(\mathcal{S}) \leq \sum_{j \in \mathcal{R}} \sum_{i \in [k]} s_{i,j}$.

PROOF. This follows from Claim 10. Therefore, \mathcal{R} gets the highest score $|\mathcal{R}|k$ items, excluding the large items. \square

CLAIM 12. Let OPT be the optimal value of MoA for a given instance. Then $OPT \leq \frac{\sum_{j \in \mathcal{R}} \sum_{i \in [k]} s_{i,j}}{|\mathcal{R}|}$.

PROOF. There is a partition of the items into m groups, such that the sum of the scores of each group is $\geq OPT$. Let $AVG_{\mathcal{R}} = \frac{\sum_{j \in \mathcal{R}} \sum_{i \in [k]} s_{i,j}}{|\mathcal{R}|}$. If $|\mathcal{R}| = m$, then the claim is obvious. Suppose for contradiction, that $OPT > AVG_{\mathcal{R}}$ when $1 \leq |\mathcal{R}| < m$. We have $m' = m - |\mathcal{R}|$ groups outside the reduced set. Therefore, we only have m' large items, that can be distributed to at most m' groups in any optimal solution. The remaining $\geq |\mathcal{R}|$ groups in the optimal solution (that do not receive any large item) would each need to get k items, summing up to $\geq OPT$. Therefore, there must exist a subset of items, say \mathcal{S} , of cardinality $|\mathcal{R}|k$, excluding the large items, such that their score sums up to $\geq |\mathcal{R}|OPT$. From Claim 11, therefore, $AVG_{\mathcal{R}}|\mathcal{R}| \geq Sum(\mathcal{S}) \geq |\mathcal{R}|OPT$. However, this contradicts the assumption that $AVG_{\mathcal{R}} < OPT$. This completes the proof. \square

Hence, we conclude from Claim 12 that $OPT \leq AVG_{\mathcal{R}}$. Suppose the minimum total score is realized by a group $r \notin \mathcal{R}$. Clearly, $s_{r,1} \geq \sum_{p \in [k-1]} s_{q,p} \forall q \in \mathcal{R}$. Since we assign the items in non-increasing order of their scores, $s_{q,k} \leq \frac{\sum_{p \in [k-1]} s_{q,p}}{k-1} \forall q \in \mathcal{R}$, hence, $s_{r,1} \geq (1 - \frac{1}{k}) \frac{\sum_{j \in \mathcal{R}, i \in [k] s_{i,j}}}{|\mathcal{R}|}$. Therefore, $s_{r,1} \geq (1 - \frac{1}{k}) OPT$. Hence, for $k \geq 2$, the realized minimum total score in this case is $\geq \frac{1}{2} OPT$.

Now, consider the case when the minimum total score is realized by some group in \mathcal{R} . Let p and q be the groups with the lowest and highest sum of scores in \mathcal{R} respectively. By definition of \mathcal{R} , there exists at least one iteration after all the groups have received one item each (that is, some iteration $> m$), when the total score of q was lower than p , and hence q got assigned an item favored over p . Let $k' \in [2, \dots, k]$ be the highest index such that the q received the k'^{th} item while p was not full. Let p have $k'' < k$ items assigned at that time. Since q was favored over p , $\sum_{i \in [k'']} s_{p,i} \geq \sum_{j \in [k'-1]} s_{q,j}$. Also, $s_{q,k'} \leq s_{p,k''}$. After this, p received $k - k''$ more items, before q received any of the items in $\{v_{q,k'+1}, \dots, v_{q,k}\}$, if $k' < k$ (otherwise, q would not receive any items after this). Clearly, $s_{q,\ell} \leq s_{p,k}$ for $\ell \in \{k' + 1, \dots, k\}$. Hence, the total score of q is $\sum_{i \in [k]} s_{q,i} \leq \sum_{i \in [k'']} s_{p,i} + s_{p,k''} + (k - k') s_{p,k}$. Therefore, $\sum_{i \in [k]} s_{q,i} \leq \sum_{i \in [k'']} s_{p,i} + s_{p,k''} + (k - k'') s_{p,k} + (k'' - k') s_{p,k} \leq S_p + s_{p,k''} + (k'' - k') s_{p,k}$. Now, $k'' - k' < k$, and because we consider items in sorted order, $s_{p,\ell} \geq s_{p,k} \forall \ell \in [k]$. Hence, $s_{p,k''} + (k'' - k') s_{p,k} \leq s_{p,k''} + \sum_{i \in \{1, \dots, k\} \setminus k''} s_{p,i} \leq S_p$. Hence, $S_q \leq 2S_p$. Since $S_q \geq AVG_{\mathcal{R}} \geq OPT$, the realized minimum total score $S_p \geq \frac{OPT}{2}$. This concludes the proof.

Algorithm 1 LEARNORDER

Parameters: Number of users n , number of groups m , group size k , confidence δ

Set δ^* as in Theorem 13

Generate a Erdos-Renyi random graph $G \sim \mathcal{G}(n, \frac{\log(n)}{n})$

Let $diam(G)$ be the diameter of G .

Let E_1, \dots, E_ℓ be a partition of the edges of G into ℓ bins got using an (approximate) minimum edge coloring G .

for $i = 1$ to ℓ **do**

Divide E_i arbitrarily into $b_i := \lceil |E_i|/m \rceil$ disjoint bins $\{B_1^i, \dots, B_{b_i}^i\}$.

for $j = 1$ to b_i **do**

Play k -partitions corresponding to B_j^i for $O\left(\frac{diam(G)^2}{\Delta^2} \ln\left(\frac{1}{\delta^*}\right)\right)$ rounds each.

Estimate $\hat{s}_p - \hat{s}_q$ for all edges $(p, q) \in B_j^i$

end for

end for

For each k , estimate \hat{s}_k by summing the estimates for $\hat{s}_i - \hat{s}_j$ along the shortest path in G from 1 to k . If no path exists for node k , set $\hat{s}_k = 0$;

Return $\hat{\sigma} = \text{argsort}(\hat{s})$.

6. LEARNING SCORE VECTOR AND GUARANTEES

In this section, we propose an algorithm for adaptively learning the optimal ordering corresponding to the score vector of a pairwise compatibility matrix. The learning proceeds in rounds and for every group \mathbf{S}_i^t in a chosen k -partition $\Pi(\mathbf{S}_1^t, \dots, \mathbf{S}_m^t)$ at round t , we assume that we receive

a iid noisy version of the happiness of the group as the response i.e. $H(\mathbf{S}_i^t) + \eta_i^t$ where $\forall t, i, \eta_i^t \in [-b, b]$ for some $b > 0$ and $\mathbb{E}(\eta_i^t) = 0$.

Our goal is to learn the ordering corresponding to the score vector $\mathbf{s} \in \mathbb{R}^n$ of the pairwise compatibility matrix $\mathbf{W} \in \mathbb{R}^{n \times n}$, by choosing groups adaptively for $T(n, k, \delta)$ rounds. Here k is the size of groups chosen in each round, and δ is the failure probability for learning a wrong ordering. Once the ordering is learned, we can compute the optimal (or approximately optimal) partition for the various objectives by sorting the scores and invoking Theorems 6, 7 and 9.

The algorithm to learn the ordering is given in Algorithm 1. The Algorithm LEARNORDER begins by generating a random Erdos-Renyi graph G where the probability of an edge being present is $\frac{\log(n)}{n}$. Thus the expected number of edges in the graph is $n \log(n)$. The edges of G are then partitioned into disjoint pieces using an approximate $O(\Sigma)$ edge coloring where Σ is the maximum degree of the G . For each of these pieces, for every edge (i, j) in the piece, groups $\{i, \mathbf{S}_{ij}\}$ and $\{j, \mathbf{S}_{ij}\}$ are chosen where \mathbf{S}_{ij} is a fixed $k - 1$ sized set that does not contain i or j . The idea is that, by obtaining the un-normalized happiness values $h_i = (s_i + \sum_{l \in \mathbf{S}_{ij}} s_l)^2 + k^2 \eta_i$ and $h_j = (s_j + \sum_{l \in \mathbf{S}_{ij}} s_l)^2 + k^2 \eta_j$ for these two groups over multiple rounds, one can compute a estimate of the difference of the corresponding scores $s_i - s_j$ with high confidence. As we only require the relative ordering to be correct, we can without loss of generality, set $\hat{s}_1 = 0$ and compute the remaining scores using the following procedure: For node k , we find the shortest path in G that connects 1 and k and sum the differences along this path (w.h.p G is connected and so there exists at least one path connecting 1 and k). Each of these differences are estimates and hence the confidence in the sum of the these estimates depend on the diameter of G .

We formally state the guarantee for LEARNORDER below.

THEOREM 13. (PAC guarantee for LEARNORDER) *Let $\mathbf{W} \in \mathbb{R}^{n \times n}$ be a score based compatibility matrix with score vector $\mathbf{s} \in \mathbb{R}^n$. Let $\Delta_{\min} = \min_{i \neq j} |s_i - s_j|$, $\Delta = 2k s_{\min} \Delta_{\min} - \Delta_{\min}^2$ and let $\delta^* = \left(1 - \left(\exp\left(-\frac{\delta}{k}\right)\right)\right)/m$. Then, algorithm LEARNORDER (Algorithm 1) outputs a permutation $\hat{\sigma}$ whose ordering is same as that of \mathbf{s} with probability at least $1 - \delta$ after $O\left(\frac{|E|}{m} \frac{diam(G)^2}{\Delta^2} \ln\left(\frac{1}{\delta^*}\right)\right)$ rounds.*

Proof Sketch: For a chosen edge pair (i, j) , let h_i and h_j denote the unnormalized happiness values obtained by choosing the groups (i, \mathbf{S}_{ij}) and (j, \mathbf{S}_{ij}) where $|\mathbf{S}_{ij}| = k - 1$ and $i, j \notin \mathbf{S}_{ij}$. We have,

$$\sqrt{h_i} - \sqrt{h_j} = \sqrt{(s_i + a)^2 + k^2 \eta_i} - \sqrt{(s_j + a)^2 + k^2 \eta_j}$$

where η_i and η_j correspond to the bounded random noise and $a = \sum_{l \in \mathbf{S}_{ij}} s_l$. If the noise η_i, η_j were not present, then the

difference $\sqrt{h_i} - \sqrt{h_j} = s_i - s_j$, which is what we want to estimate. Nevertheless, we can control the noise by playing this pair of groups for $O\left(\frac{diam(G)^2}{\Delta^2} \ln\left(\frac{1}{\delta^*}\right)\right)$ rounds as in the Algorithm and averaging the happiness values obtained to obtain estimates \hat{s}_{ij} . In this case, after these many rounds, we have with probability at least $1 - \delta^*$,

$$\frac{(s_i - s_j) - \Delta_{\min}}{diam(G)} \leq \hat{s}_{ij} \leq \frac{(s_i - s_j) + \Delta_{\min}}{diam(G)} \quad \forall i, j$$




Figure 1: Normalized error between the estimated and true weights (y-axis) of LEARNORDER. The left plot is for a random graph and the right is for a Facebook graph instance.

When computing the estimate of a pair (s_1, s_k) not in the edge set E , the algorithm sums up the estimated weights on the shortest path from 1 to k . As the shortest path is at most $diam(G)$ long by definition, we obtain estimates for all pairs of the form (s_1, s_k) such that all the estimated values \hat{s}_{1k} satisfies with probability $1 - \delta^*$, $|\hat{s}_{1k} - s_{1k}| \leq \Delta_{\min} \forall k$.

Thus under the above condition, if we fix $\hat{s}_1 = 0$ and obtain values for all other vertices, we can sort them to produce an ordering. It is easy to see that this ordering will exactly correspond to the ordering of the actual score vector \mathbf{s} .

Remark: The sample complexity (i.e., the number of groups to be chosen) by the above Theorem depends on the diameter of the random graph G . It is known that for large enough n , $diam(G)$ is concentrated sharply around $2 \frac{\log(n)}{\log(n/2)}$ and the number of edges behaves as $O(n \log(n))$.

7. EXPERIMENTS

We assess the quality of the LEARNORDER algorithm using simulated and real data. The graph instances we chose were the following. (a) *Random*: for this synthetically generated dataset, we fixed the group size $k = 4$ and the graph size $n = 16$. We added uniform noise between $[-1, 1]$ to the feedback in each round and the score for each item was drawn uniformly at random from $\{1, \dots, 10\}$. And (b) *Social network*: a 16 node instance was sampled from the Facebook friendship graph, built from an existing anonymized dataset of Facebook users' data (Leskovec and Krevl 2014). The dataset has 4039 user nodes and 88234 unweighted edges. We used the *Jaccard similarity coefficient* of features such as education, hometown, language, etc to obtain scores for the users. The performance of LEARNORDER is shown in Figure 1 (averaged over 30 runs), and is in terms of the normalized error between the estimated score vector and the true score vector. It decreases as the number of rounds t increases.

In addition to showing that the learning algorithm indeed converges, the experiments add empirical support to the fact that the number of rounds needed to learn the true scores within 10% normalized error is very practical (for instance 70 rounds for the Facebook subgraph of size 16). Since the

weight matrix was filled using a pairwise similarity measure (using demographical and other user specific metadata) and no intrinsic score was assumed, the experiment shows that the weight matrix is naturally low rank (allowing us to learn the scores very well) for this dataset. As a consequence, for this dataset, we could infer that the users' intrinsic characteristics reasonably determine who they are friends with.

8. CONCLUSIONS

We studied the problem of grouping a set of users using their pairwise compatibilities. We first showed hardness and inapproximability results when no assumptions on the pairwise compatibility values are made. We then studied the intrinsic score model for the compatibility, a model that is not only simple and straight forward but also very similar to the popular Bradley–Terry–Luce (BTL) model for pairwise comparisons. Under this model, we related the optimal groupings to homophilous and heterophilous groupings which are well studied in the psychology literature. We proposed the LEARNORDER algorithm, which after choosing a small number of groups, adaptively learns the best ordering corresponding to the score vector of the pairwise compatibility matrix. Our experiments on both synthetic and real datasets demonstrate the efficacy of our algorithm.

We note that there may be several applications where the pairwise compatibilities between users/items may depend on multiple features (instead of one) and the pairwise compatibility matrix can in general be low rank (instead of being score based). In such cases, our framework can be slightly modified to incorporate a matrix completion subroutine to recover the low rank compatibility matrix. However, the results regarding the optimality of the homophilous/heterophilous partitions do not follow. The analysis of this is beyond the scope of the current work.

In the future, we would like to consider other relevant structures for the happiness index and develop algorithms for the same, possibly with statistical as well as computational guarantees.

REFERENCES

- [1] M. I. Alpert and W. T. Anderson. Optimal heterophily and communication effectiveness: Some empirical findings. *Journal of Communication*, 23(3):328–343, 1973.
- [2] M. Baccara and L. Yariv. Homophily in peer groups. Available at SSRN 2045156, 2012.
- [3] N. Bansal, U. Feige, R. Krauthgamer, K. Makarychev, V. Nagarajan, J. Naor, and R. Schwartz. Min-max graph partitioning and small set expansion. In *FOCS, 2011*, pages 17–26. IEEE, 2011.
- [4] F. Bistaffa, A. Farinelli, and S. D. Ramchurn. Sharing rides with friends: a coalition formation algorithm for ridesharing. *AAAI*, 2015.
- [5] R. A. Bradley and M. E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.
- [6] J. Brindley, L. M. Blaschke, and C. Walti. Creating effective collaborative learning groups in an online environment. *The International Review of Research in Open and Distributed Learning*, 10(3), 2009.
- [7] R. Cordone and F. Maffioli. On the complexity of graph tree partition problems. *Discrete Applied Mathematics*, 134(1):51–65, 2004.
- [8] R. Dies. Research on group psychotherapy: Overview and clinical applications. *Group therapy in clinical practice*, pages 473–518, 1993.
- [9] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965.
- [10] B. Elbaum, S. Vaughn, M. Hughes, and S. Moody. Grouping practices and reading outcomes for students with disabilities. *Exceptional children*, 65(3):399, 1999.
- [11] D. Eppstein, J. Feigenbaum, and C. Li. Equipartitions of graphs. *Discrete mathematics*, 91(3):239–248, 1991.
- [12] T. Feder, P. Hell, S. Klein, and R. Motwani. Complexity of graph partition problems. In *Proceedings of the thirty-first annual ACM STOC*, pages 464–472. ACM, 1999.
- [13] B. R. Foorman and J. Torgesen. Critical elements of classroom and small-group instruction promote reading success in all children. *Learning Disabilities Research & Practice*, 16(4):203–212, 2001.
- [14] M. R. Garey and D. S. Johnson. A guide to the theory of np-completeness. *WH Freeman*, 1979.
- [15] D. R. Gaur, R. Krishnamurti, and R. Kohli. The capacitated max k-cut problem. *Mathematical Programming*, 115(1):65–72, 2008.
- [16] M. Grötschel and Y. Wakabayashi. Facets of the clique partitioning polytope. *Mathematical Programming*, 47(1-3):367–387, 1990.
- [17] Y. He, Z. Tan, J. Zhu, and E. Yao. κ -partitioning problems for maximizing the minimum load. *Computers & Mathematics with Applications*, 46(10):1671–1681, 2003.
- [18] S. K. Horwitz. The compositional impact of team diversity on performance: Theoretical considerations. *Human resource development review*, 4(2):219–245, 2005.
- [19] S. K. Horwitz and I. B. Horwitz. The effects of team diversity on team outcomes: A meta-analytic review of team demography. *Journal of management*, 33(6):987–1015, 2007.
- [20] P. F. Lazarsfeld and R. K. Merton. Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society*, 18(1):18–66, 1954.
- [21] T. Liang, C. Liu, T.-M. Lin, and B. Lin. Effect of team diversity on software project performance. *Industrial Management & Data Systems*, 107(5):636–653, 2007.
- [22] R. D. Luce. Individual choice behavior: A theoretical analysis. *New York*, 115:191–243, 1959.
- [23] M. McPherson, L. Smith-Lovin, and J. M. Cook. Birds of a feather: Homophily in social networks. *Annual review of sociology*, pages 415–444, 2001.
- [24] A. Singla, E. Horvitz, P. Kohli, and A. Krause. Learning to hire teams. *arXiv:1508.02823*, 2015.
- [25] G. K. Stahl, K. Mäkelä, L. Zander, and M. L. Maznevski. A look at the bright side of multicultural team diversity. *Scandinavian Journal of Management*, 26(4):439–447, 2010.
- [26] Z. Svitkina and É. Tardos. Min-max multiway cut. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 207–218. Springer, 2004.
- [27] J. R. Terborg, C. Castore, and J. A. DeNinno. A longitudinal field investigation of the impact of group composition on group performance and cohesion. *J. Personality and Social Psychology*, 34(5):782, 1976.
- [28] R. Watrigant, M. Bougeret, R. Giroudeau, and J. König. On the sum-max graph partitioning problem. *Theoretical Computer Science*, 540:143–155, 2014.
- [29] S. Zhu, D. Wang, and T. Li. Data clustering with size constraints. *Knowledge-Based Systems*, 23(8):883–889, 2010.