

Sync-Rank: Robust Ranking, Constrained Ranking and Rank Aggregation via Eigenvector and SDP Synchronization

Mihai Cucuringu
E-mail: mihai@math.ucla.edu

Abstract—We consider the classical problem of establishing a statistical ranking of a set of n items given a set of inconsistent and incomplete pairwise comparisons between such items. Instantiations of this problem occur in numerous applications in data analysis (e.g., ranking teams in sports data), computer vision, and machine learning. We formulate the above problem of ranking with incomplete noisy information as an instance of the *group synchronization* problem over the group $SO(2)$ of planar rotations, whose usefulness has been demonstrated in numerous applications in recent years in computer vision and graphics, sensor network localization and structural biology. Its least squares solution can be approximated by either a spectral or a semidefinite programming (SDP) relaxation, followed by a rounding procedure. We perform extensive numerical simulations on both synthetic and real-world data sets, which show that our proposed method compares favorably to other ranking methods from the recent literature. Existing theoretical guarantees on the group synchronization problem imply lower bounds on the largest amount of noise permissible in the data while still achieving an approximate recovery of the ground truth ranking. We propose a similar synchronization-based algorithm for the rank-aggregation problem, which integrates in a globally consistent ranking many pairwise rank-offsets or partial rankings, given by different rating systems on the same set of items, an approach which yields significantly more accurate results than other aggregation methods, including *Rank-Centrality*, a recent state-of-the-art algorithm. Furthermore, we discuss the problem of semi-supervised ranking when there is available information on the ground truth rank of a subset of players, and propose an algorithm based on SDP which is able to recover the ranking of the remaining players, subject to such hard constraints. Finally, synchronization-based ranking, combined with a spectral technique for the densest subgraph problem, makes it possible to extract partial rankings that other methods are not able to find, in other words, to identify the rank of a small subset of players whose pairwise rank comparisons are less noisy than the rest of the data. We discuss a number of related open questions which we defer for future investigation.

Index Terms—ranking, angular synchronization, spectral algorithms, semidefinite programming, rank aggregation, partial rankings, least squares, singular value decomposition, densest subgraph problem.

1 INTRODUCTION

THE problem of ranking a set of n players from incomplete and noisy cardinal or ordinal pairwise measurements has received a great deal of attention in the recent literature. Most often, the data is incomplete, especially when n is large, but also very noisy, with a large fraction of the pairwise measurements being both incorrect and inconsistent with respect to the existence of an underlying total ranking. The goal is to recover a global ranking of the players that is consistent with the given data as best as possible.

The analysis of many modern large-scale data sets implicitly requires various forms of ranking to allow for the identification of the most important entries or features. A myriad of such instances appear throughout various disciplines, including internet-related applications such as the famous Google search engine [29], eBay’s feedback-based reputation mechanism [37], Amazon’s Mechanical Turk system, or the Netflix recommendation system [10].

Traditional ranking methods, most often coming from the social choice theory literature, have proven less efficient in dealing with nowadays data. Most of this literature has

been developed with ordinal comparisons in mind, while much of the current data deals with cardinal (numerical) scores for the pairwise comparisons, such as sports data, where the outcome of a match is often a score.

There exists a very rich literature on ranking, and it would be impossible to give a fair account of it here. Instead, we will mention some of the methods that relate to our work, and especially methods we compare against. A very common approach in the ranking literature is to treat the input rankings as data generated from a probabilistic model, and then learn the Maximum Likelihood Estimator of the input data, an idea which has been explored in depth in the machine learning community, and includes the Bradley-Terry-Luce model (BTL) [9] or the Plackett-Luce (PL) model [26]. Soufiani et al. propose an efficient Generalized Method-of-Moments algorithm for computing parameters of the PL model, by breaking the full rankings into pairwise comparisons, and then computing the parameters that satisfy a set of generalized moment conditions [3]. The recent work of [35] for ranking from a random sample of ordinal comparisons, also accounts for whether one seeks an approximately uniform quality across the ranking, or more accuracy near the top or the bottom.

The general idea of angular embedding, which we exploit in this paper, is not new, and aside from recent work by

• M. Cucuringu is with the Department of Mathematics, UCLA, 520 Portola Plaza, Mathematical Sciences Building 6363, Los Angeles, CA 90095-1555.

Singer in the context of the angular synchronization problem (ASP) [30], has also been explored by Yu [38], who observed that embedding in the angular space is significantly more robust to outliers when compared to embedding in the usual linear space. In the latter case, the traditional least squares formulation, or its l_1 norm formulation, cannot match the performance of the angular embedding approach, thus suggesting that the key to overcoming noise comes not with imposing additional constraints on the solution, but by penalizing the inconsistencies between the measurements. Yu’s proposed spectral method returns very satisfactory results in terms of robustness to noise for an image reconstruction problem. Recently, Braxton et al. [28] propose an l_1 -norm formulation for ranking with cardinal data, as an alternative to a previous l_2 -norm formulation [22].

Contribution of our paper. The contribution of our proposed Sync-Rank algorithm can be summarized as follows.

- We make an explicit connection between ranking and the angular synchronization problem, and use existing spectral and SDP relaxations for the latter problem to compute robust global rankings.
- We perform a very extensive set of numerical simulations comparing our proposed method with existing state-of-the-art algorithms from the ranking literature, across a variety of synthetic measurement graphs and noise models, both for numerical (*cardinal*) and binary (*ordinal*) pairwise comparisons. In addition, we compare the algorithms on two real data sets: scores of soccer games in the English Premier League, and NCAA College Basketball games. Overall, we compare (in most instances, favorably) to the two recently proposed state-of-the-art algorithms, *Serial-Rank* [18], and *Rank-Centrality* [27].
- Furthermore, we propose and compare to a very simple ranking method based on Singular Value Decomposition (that we currently investigate in a separate ongoing work¹), which may be of independent interest as its performance is comparable to that of the recent *Serial-Rank* algorithm, and superior to it in the setting where only the magnitude of the rank offset is available, but not its sign.
- We propose a method for ranking in the semi-supervised setting where a subset of the players have a prescribed rank to be enforced as a hard constraint.
- We show that our approach is applicable to the rank aggregation problem of integrating ranking information from multiple rating systems that provide independent, incomplete and inconsistent pairwise comparisons for the same set of players.
- Finally, we show that by combining Sync-Rank with recent algorithms for the *densest subgraph problem*, we are able to identify *planted partial rankings*, which other methods are not able to extract.

The advantage of Sync-Rank stems from the fact that it is a computationally simple, algorithm that is model independent and relies exclusively on the available (ordinal or cardinal) data. Existing theoretical guarantees from the recent literature on the group synchronization problem

[4], [30] trivially translate to lower bounds for the largest amount of noise permissible in the measurements that still allows for an exact or almost exact recovery.²

The remainder of this paper is organized as follows. Section 2 summarizes related methods against which we compare. Section 3 is a review of the angular synchronization problem and existing results from the literature. Section 4 describes the Sync-Rank algorithm for ranking via eigenvector and SDP-based synchronization. Section 5 provides an extensive numerical comparison of Sync-Rank with other recent algorithms, across a variety of synthetic and real data sets. Section 6 considers the rank aggregation problem, which we solve efficiently via the same spectral and SDP relaxations of the angular synchronization problem. In Section 7 we consider the constrained ranking problem and propose a modified SDP-based synchronization algorithm for it. Section 8 summarizes several open problems related to ranking, on which we expand upon in Appendix D. Section 9 is a summary and discussion. In Appendix A, we propose an algorithm for extracting partial rankings from comparison data. Appendix B summarizes the recent *Serial-Rank* algorithm. Appendix C summarizes the *Rank-Centrality* algorithm, and our proposed version for a single rating system. In Appendix E we show recovered rankings for the English Premier League 2013-2014 season.

2 RELATED METHODS

2.1 Serial Rank and Generalized Linear Models

In very recent work [18], Fogel et al. propose a seriation algorithm for ranking a set of players given noisy incomplete pairwise comparisons. The gist of their approach is to assign similar rankings to players that compare similarly with all other players. They do so by constructing a similarity matrix from the available pairwise comparisons, relying on existing seriation methods to reorder the similarity matrix and thus recover the final rankings. The authors make an explicit connection between the ranking problem and another related classical ordering problem, namely *seriation*, where one is given a similarity matrix between a set of n items and assumes that the items have an underlying ordering on the line such that the similarity between items decreases with their distance. In other words, the more similar two items are, the closer they should be in the proposed solution. By and large, the goal of the seriation problem is to recover the underlying linear ordering based on unsorted, inconsistent and incomplete pairwise similarity information. We briefly summarize their approach in Appendix B, and note that it closely resembles our proposed SVD-Rank algorithm.

2.2 The Rank-Centrality algorithm

Negahban et al. [27] recently proposed an iterative algorithm for the rank aggregation problem of integrating ranking information from multiple ranking systems, by estimating scores for the items from the stationary distribution of a certain random walk on the graph of items, where each edge encodes the outcome of pairwise comparisons. We summarize their approach in Appendix C, and compare

1. With Alexandre d’Aspremont and Fajwel Fogel, INRIA and École Normale Supérieure, Paris.

2. We point out that a perfect recovery of the angles in the ASP is not a necessary condition for a perfect recovery of the underlying ground truth ranking, since it suffices that only the relative ordering of the angles is preserved.

against it in the setting of the rank aggregation problem. However, for the case of a single rating system, we propose and compare to a variant of their algorithm, which we also detail in the same Appendix C.

2.3 Ranking via Singular Value Decomposition

An additional ranking method we propose, and compare against, is based on the traditional Singular Value Decomposition (SVD) method. The applicability of the SVD-Rank approach stems from the observation that, in the case of cardinal measurements (25), the noiseless matrix of rank offsets C is a skew-symmetric matrix of even rank 2 since

$$R = re^T - er^T \quad (1)$$

where e denotes the all-ones column vector. In the noisy case, C is a random perturbation of a rank-2 matrix. We consider the top two singular vectors of C , order their entries by their size, extract the resulting rankings, and choose between the first and second singular vector based on whichever one minimizes the number of upsets. Note that since the singular vectors are obtained up to a global sign, we choose the ordering which minimizes the number of upsets. Though a rather naive approach, SVD-Rank returns, under the multiplicative uniform noise model, results that are comparable to those of very recent algorithms such as Serial-Rank [18] and Rank-Centrality [27]. A previous application of SVD to ranking has been explored in Gleich and Zhukov [20], for studying relational data as well as developing a method for interactive refinement of the search results. To the best of our knowledge, we are not aware of other work that considers SVD-based ranking for the setting considered in this paper. An interesting research direction we discuss in Appendix D.1, is to analyze the performance of SVD-Rank using tools from the random matrix theory literature on rank-2 deformations of random matrices.

The SVD-Rank approach is similar to that of the Serial-Rank algorithm detailed in Appendix B, since the singular vectors of C are nothing but eigenvectors of CC^T , the only differences being the addition of a multiple of the all ones matrix as in (42) and working with the Fiedler vector of the resulting Laplacian matrix. Finally, we mention that in ongoing work¹ we use SVD-Rank in the context of the *seriation* problem, where one aims to recover the global ordering of the items (up to a global sign) using only information on the absolute value of the rank offsets (see Section D.3 for additional details), and note that SVD-Rank performs significantly better than its Serial-Rank counterpart.

2.4 Ranking via Least Squares

We also compare our proposed ranking method with the more traditional least-squares approach. Assuming the number of edges in G is given by $m = |E(G)|$, we denote by B the edge-vertex incidence matrix of size $m \times n$

$$B_{ij} = \begin{cases} 1 & \text{if } (i, j) \in E(G), \text{ and } i > j \\ -1 & \text{if } (i, j) \in E(G), \text{ and } i < j \\ 0 & \text{if } (i, j) \notin E(G), \end{cases} \quad (2)$$

and by y the vector of length $m \times 1$ which contains the pairwise rank measurements $y(e) = C_{ij}$, for all edges $e =$

$(i, j) \in E(G)$. We obtain the least-squares solution to the ranking problem by solving the following minimization

$$\underset{x \in \mathbb{R}^n}{\text{minimize}} \quad \|Bx - y\|_2^2. \quad (3)$$

We point out here the work of Hirani et al. [21], who show that the problem of least-squares ranking on graphs has far-reaching rich connections with various other research areas, including spectral graph theory and multilevel methods for graph Laplacian systems, and Hodge decomposition theory.

3 THE GROUP SYNCHRONIZATION PROBLEM

Finding group elements from noisy measurements of their ratios is known as the *group synchronization* problem. For example, the synchronization problem over the special orthogonal group $SO(d)$ consists of estimating a set of n unknown $d \times d$ matrices $R_1, \dots, R_n \in SO(d)$ from noisy measurements R_{ij} of a subset of their pairwise ratios $R_i^{-1}R_j$. The least squares solution to synchronization aims to minimize the sum of squared deviations

$$\underset{R_1, \dots, R_n \in SO(d)}{\text{minimize}} \quad \sum_{(i,j) \in E} w_{ij} \|R_i^{-1}R_j - R_{ij}\|_F^2, \quad (4)$$

where $\|\cdot\|$ denotes the Frobenius norm, and w_{ij} are non-negative weights representing the confidence in the noisy pairwise measurements R_{ij} . Spectral and semidefinite programming (SDP) relaxations for solving an instance of the above synchronization problem were originally introduced and analyzed by Singer [30] in the context of angular synchronization, over the group $SO(2)$ of planar rotations, where one is asked to estimate n unknown angles

$$\theta_1, \dots, \theta_n \in [0, 2\pi), \quad (5)$$

given m noisy measurements Θ_{ij} of their offsets

$$\Theta_{ij} = \theta_i - \theta_j \pmod{2\pi}. \quad (6)$$

The difficulty of the problem is amplified on one hand by the amount of noise in the offset measurements, and on the other hand by the fact that $m \ll \binom{n}{2}$, i.e., only a very small subset of all possible pairwise offsets are measured. In general, one may consider other groups \mathcal{G} (such as $SO(d)$, $O(d)$) for which there are available noisy measurements g_{ij} of ratios between the group elements $g_{ij} = g_i g_j^{-1}$, $g_i, g_j \in \mathcal{G}$. The set E of pairs (i, j) for which a ratio of group elements is available can be realized as the edge set of a graph $G = (V, E)$, $|V| = n$, $|E| = m$, with vertices corresponding to the group elements g_1, \dots, g_n , and edges to the available pairwise measurements $g_{ij} = g_i g_j^{-1}$.

In [30], Singer analyzed the following noise model, where each edge in the measurement graph G is present with probability p , and each available measurement is either correct with probability $1 - \eta$ or a random measurement with probability η . For such a noise model with outliers, the available measurement matrix Θ is given by the mixture

$$\Theta_{ij} = \begin{cases} \theta_i - \theta_j & \text{for a correct edge, w.p. } p(1 - \eta) \\ \sim \mathcal{U}(S^1) & \text{for an incorrect edge, w.p. } p\eta \\ 0 & \text{for a missing edge, w.p. } 1 - p. \end{cases} \quad (7)$$

where $\mathcal{U}(S^1)$ denotes the uniform distribution over the unit circle. Using tools from random matrix theory, in particular

rank-1 deformations of large random matrices [17], Singer showed in [30] that for the complete graph ($p = 1$), the spectral relaxation for angular synchronization given by (12) and summarized in the next Section 3.1, undergoes a phase transition phenomenon, with the top eigenvector of the Hermitian matrix H in (9) exhibiting above random correlations with the ground truth as soon as

$$1 - \eta > \frac{1}{\sqrt{n}}. \quad (8)$$

In other words, even for very small values of $1 - \eta$ (thus a large noise level), the eigenvector synchronization method yields a solution which (prior to the normalization in (10)), is shown to exhibit above random correlations with the ground truth angles if there are enough pairwise measurements available, i.e., whenever $n(1 - \eta)^2$ is large enough. For the general case of Erdős-Rényi graphs the above threshold probability can be extended to $O\left(\frac{1}{\sqrt{np}}\right)$. We remark that while there exist no theoretical guarantees on the exactness of the eigenvector method for angular synchronization, in practice this relaxation has been observed to be nearly tight for small enough levels of noise η . On the other hand, for the SDP relaxation of the Least Unsquared Deviation (LUD) formulation [33], there do exist exact recovery guarantees under the above Erdős-Rényi graph measurement model.

3.1 Spectral Relaxation

Following the approach introduced in [30], we build the $n \times n$ sparse Hermitian matrix $H = (H_{ij})$ whose elements are either zero or points on the unit circle in the complex plane

$$H_{ij} = \begin{cases} e^{i\theta_{ij}} & \text{if } (i, j) \in E \\ 0 & \text{if } (i, j) \notin E. \end{cases} \quad (9)$$

In an attempt to preserve the angle offsets as best as possible, Singer considers the following maximization problem

$$\underset{\theta_1, \dots, \theta_n \in [0, 2\pi)}{\text{maximize}} \sum_{i,j=1}^n e^{-i\theta_i} H_{ij} e^{i\theta_j} \quad (10)$$

which gets incremented by +1 whenever an assignment of angles θ_i and θ_j perfectly satisfies the given edge constraint $\theta_{ij} = \theta_i - \theta_j \bmod 2\pi$ (i.e., for a *good* edge), while the contribution of an incorrect assignment (i.e., of a *bad* edge) will be uniformly distributed on the unit circle in the complex plane. Note that (10) is equivalent to the formulation in (4) by exploiting properties of the Frobenius norm, and relying on the fact that it is possible to represent group elements for the special case of $\text{SO}(2)$ as complex-valued numbers. Since the non-convex optimization problem in (10) is difficult to solve computationally [39], Singer introduced the following spectral relaxation

$$\underset{z_1, \dots, z_n \in \mathbb{C}; \sum_{i=1}^n |z_i|^2 = n}{\text{maximize}} \sum_{i,j=1}^n \bar{z}_i H_{ij} z_j \quad (11)$$

by replacing the individual constraints $z_i = e^{i\theta_i}$ having unit magnitude by the much weaker single constraint $\sum_{i=1}^n |z_i|^2 = n$. Next, we recognize the resulting maximization problem in (11) as the maximization of a quadratic form whose solution is given by the top eigenvector of the Hermitian matrix H , which has an orthonormal basis over

\mathbb{C}^n , with real eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$ and corresponding eigenvectors v_1, v_2, \dots, v_n . In other words, the spectral relaxation of the non-convex optimization problem in (10) is given by

$$\underset{\|z\|^2 = n}{\text{maximize}} z^* H z \quad (12)$$

which can be solved via a simple eigenvector computation, by setting $z = v_1$, where v_1 is the top eigenvector of H , satisfying $Hv_1 = \lambda_1 v_1$, with $\|v_1\|^2 = n$, corresponding to the largest eigenvalue λ_1 . Before extracting the final estimated angles, we consider the following normalization of H using the diagonal matrix D , whose diagonal elements are given by $D_{ii} = \sum_{j=1}^n |H_{ij}|$, and define

$$\mathcal{H} = D^{-1} H, \quad (13)$$

which is similar to the Hermitian matrix $D^{-1/2} H D^{-1/2}$

$$\mathcal{H} = D^{-1/2} (D^{-1/2} H D^{-1/2}) D^{1/2}.$$

Thus, \mathcal{H} has n real eigenvalues $\lambda_1^{\mathcal{H}} > \lambda_2^{\mathcal{H}} \geq \dots \geq \lambda_n^{\mathcal{H}}$ with corresponding n orthogonal (complex valued) eigenvectors $v_1^{\mathcal{H}}, \dots, v_n^{\mathcal{H}}$, with $\mathcal{H}v_i^{\mathcal{H}} = \lambda_i^{\mathcal{H}} v_i^{\mathcal{H}}$. We define the estimated rotation angles $\hat{\theta}_1, \dots, \hat{\theta}_n$ using the top eigenvector $v_1^{\mathcal{H}}$ via

$$e^{i\hat{\theta}_i} = \frac{v_1^{\mathcal{H}}(i)}{|v_1^{\mathcal{H}}(i)|}, \quad i = 1, 2, \dots, n \quad (14)$$

Note that the estimation of the rotation angles $\theta_1, \dots, \theta_n$ is up to an additive phase since $e^{i\alpha} v_1^{\mathcal{H}}$ is also an eigenvector of \mathcal{H} for any $\alpha \in \mathbb{R}$. We point out that the only difference between the above approach and the angular synchronization algorithm in [30] is the normalization (13) of the matrix prior to the computation of the top eigenvector, considered in our previous work [12], and formalized in [31] via the notion of *graph connection Laplacian* $\mathcal{L} = D - H$.

3.2 Semidefinite Programming (SDP) Relaxation

A second relaxation proposed in [30] as an alternative to the spectral relaxation, is via the following SDP formulation. In an attempt to preserve the angle offsets as best as possible, one may consider the following maximization problem

$$\sum_{i,j=1}^n e^{-i\theta_i} H_{ij} e^{i\theta_j} = \text{trace}(H\Upsilon), \quad (15)$$

where Υ is the unknown $n \times n$ rank-1 Hermitian matrix

$$\Upsilon_{ij} = e^{i(\theta_i - \theta_j)} \quad (16)$$

with ones in the diagonal $\Upsilon_{ii}, \forall i = 1, 2, \dots, n$. Note that, with the exception of the rank-1 constraint, all the remaining constraints are convex and altogether define the following SDP relaxation for the angular synchronization problem

$$\begin{aligned} & \underset{\Upsilon \in \mathbb{C}^{n \times n}}{\text{maximize}} && \text{trace}(H\Upsilon) \\ & \text{subject to} && \Upsilon_{ii} = 1 \quad i = 1, \dots, n \\ & && \Upsilon \succeq 0, \end{aligned} \quad (17)$$

which can be solved via standard methods from the convex optimization literature [32]. We remark that, from a computational perspective, solving such SDP problems is computationally feasible only for relative small-sized problem (typically with several thousand unknowns, up to

about $n = 10,000$), though there exist distributed methods for solving such convex optimization problems, such as the popular Alternating Direction Method of Multipliers (ADMM) [8] which can handle large-scale problems. Alternatively, block coordinate descent methods have been shown to be extremely efficient at solving a broad class of semidefinite programs, including Max-Cut SDP relaxation and minimum nuclear norm matrix completion [36]. Finally, the very recent work of Boumal [7] also addresses the issue of solving large-scale SDP problems via a Riemannian optimization approach.

As pointed out in [30], this program is very similar to the well-known Goemans-Williamson SDP relaxation for the famous MAX-CUT problem of finding the maximum cut in a weighted graph, the only difference being the fact that here we optimize over the cone of complex-valued Hermitian positive semidefinite matrices, not just real symmetric matrices. Since the recovered solution is not necessarily of rank-1, the estimator is obtained from the best rank-1 approximation of the optimal solution matrix Θ via a Cholesky decomposition. We plot in Figure 3 the recovered ranks of the SDP relaxation for the ranking problem, and point out the interesting phenomenon that, even for noisy data, under favorable noise regimes, the SDP program still returns a rank-1 solution. The tightness of this relaxation has been explained only recently in the work of Bandeira et al. [4]. The advantage of the SDP relaxation is that it explicitly imposes the unit magnitude constraint for $e^{i\theta_i}$, which we cannot otherwise enforce in the spectral relaxation.

4 SYNC-RANK: RANKING VIA SYNCHRONIZATION

We now consider the application of the angular synchronization framework [30] to the ranking problem. The underlying idea has also been considered by Yu in the context of image denoising [38], who suggested, similar to [30], to perform the denoising step in the angular embedding space as opposed to the linear space, and observed increased robustness against sparse outliers in the measurements.

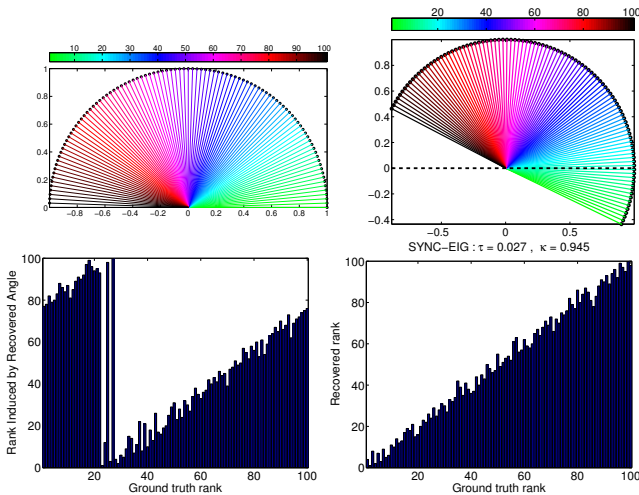


Fig. 1. Top: equidistant mapping of the rankings $1, 2, \dots, n$ in the upper half circle, for $n = 100$, where the rank of the i^{th} player is i (left), and the recovered solution at some random rotation (right). Bottom: the ranking induced by the θ angles from the ASP (left), and the ranking obtained after finding the optimal circular permutation (right).

Denote the true ranking of the n players by $r_1 < r_2 < \dots < r_n$, and assume without loss of generality that $r_i = i$,

i.e., the rank of the i^{th} player is i . In the ideal case, the ranks can be imagined to lie on a one-dimensional line, sorted from 1 to n , with the pairwise rank comparisons given, in the noiseless case, by $C_{ij} = r_i - r_j$ (for cardinal measurements) or $C_{ij} = \text{sign}(r_i - r_j)$ (for ordinal measurements). In the angular embedding space, we consider the ranks of the players mapped to the unit circle, say fixing r_1 to have a zero angle with the x -axis, and the last player r_n corresponding to an angle equal to π . In other words, we imagine the n players wrapped around a fraction of the circle, interpret the available rank-offset measurements as angle-offsets in the angular space, and thus arrive at the setup of the angular synchronization problem from Section 3.

We also remark that the modulus used to wrap the players around the circle plays an important role in the recovery process. If we choose to map the players across the entire circle, this would cause ambiguity at the end points, since the very highly ranked players will be positioned very close (or perhaps even mixed) with the very poorly ranked players. To avoid the confusion, we simply choose to map³ the n players to the upper half of the unit circle $[0, \pi]$.

In the ideal setting, the angles obtained via synchronization would be as shown in the left plot of Figure 1, from where one can easily infer the ranking of the players by traversing the upper half circle in an anti-clockwise direction. However, since the solution to the angular synchronization problem is computed up to a global shift (see for example the top right and the bottom left plots in Figure 1), an additional post-processing step is required to accurately extract the ordering of the players that best matches the given data. To this end, we simply compute the best circular permutation of the initial rankings obtained from synchronization, that minimizes the number of upsets in the given data. We illustrate this step with a noisy instance of Sync-Rank in the bottom of Figure 1, where the left plot shows the rankings induced by the initial angles recovered from synchronization, while the right one shows the final ranking, after shifting by the best circular permutation.

Denote by $s = [s_1, s_2, \dots, s_n]$ the ordering induced by the angles recovered from angular synchronization, when sorting the angles from smallest to largest, where s_i denotes the label of the player ranked on the i^{th} position. For example, s_1 denotes the player with the smallest corresponding angle θ_{s_1} . To measure the accuracy of each candidate circular permutation σ , we first compute the pairwise rank offsets associated to the induced ranking, via

$$P_\sigma(s) = (\sigma(s) \otimes \mathbf{1} - \mathbf{1} \otimes \sigma(s)) \circ A \quad (18)$$

where \otimes denotes the outer product of two vectors $x \otimes y = xy^T$, \circ denotes the Hadamard product of two matrices (entrywise product), and A is the adjacency matrix of the graph G . In other words, for an edge $(u, v) \in E(G)$, it holds that $(P_\sigma(s))_{uv} = \sigma(s)_u - \sigma(s)_v$, i.e., the resulting rank offset

3. The spectral and SDP synchronization methods could be applied in the case of non-compact groups only after “compactifying” them. In the present approach, the line is “compactified” by simply mapping it to the unit circle (or part of it), making the approach amenable to the synchronization methods and less sensitive to outliers. On the other hand, the extension to non-compact groups, such as $\text{Euc}(d)$ the Euclidean group of rigid transformation, or the general linear group of invertible matrices $\text{GL}_n(R)$ poses significant challenges.

after applying the cyclic shift⁴. Next, we choose the circular permutation minimizing the l_1 norm⁵ of the residual matrix

$$\sigma = \arg \min_{\sigma_1, \dots, \sigma_n} \frac{1}{2} \|\text{sign}(P_{\sigma_i}(s)) - \text{sign}(C)\|_1 \quad (19)$$

which counts the total number of upsets. Note that an alternative to the above error is given by

$$\sigma = \arg \min_{\sigma_1, \dots, \sigma_n} \|P_{\sigma_i}(s) - C\|_1 \quad (20)$$

which takes into account the actual magnitudes of the offsets, not just their sign. We summarize the above steps of Sync-Rank in Algorithm 1, and note that throughout the paper, we denote by SYNC (or SYNC-EIG) the version of synchronization-based ranking which relies on the spectral relaxation of the synchronization problem, and by SYNC-SDP the one obtained via the SDP relaxation. Figure 2

Algorithm 1 Summary of the Synchronization-Ranking (Sync-Rank) Algorithm

Require: $G = (V, E)$ the graph of pairwise comparisons and C the $n \times n$ matrix of pairwise comparisons (rank offsets), such that whenever $(i, j) \in E(G)$ we have available a (perhaps noisy) comparison between players i and j , either a cardinal comparison ($C_{ij} \in [-(n-1), (n-1)]$) or an ordinal comparison $C_{ij} = \pm 1$.

- 1: Map all rank offsets C_{ij} to an angle $\Theta_{ij} \in [0, 2\pi\delta)$ with $\delta \in [0, 1)$, using the transformation

$$C_{ij} \mapsto \Theta_{ij} := 2\pi\delta \frac{C_{ij}}{n-1} \quad (21)$$

We choose $\delta = \frac{1}{2}$, and hence $\Theta_{ij} := \pi \frac{C_{ij}}{n-1}$.

- 2: Build the $n \times n$ Hermitian matrix H with $H_{ij} = e^{i\Theta_{ij}}$, if $(i, j) \in E$, and $H_{ij} = 0$ otherwise, as in (9).
- 3: Solve the angular synchronization problem via either its spectral (12) or SDP (17) relaxation, and denote the recovered solution by $\hat{r}_i = e^{i\hat{\theta}_i} = \frac{v_1^R(i)}{|v_1^R(i)|}$, $i = 1, 2, \dots, n$, where v_1 denotes the recovered eigenvector
- 4: Extract the corresponding set of angles $\hat{\theta}_1, \dots, \hat{\theta}_n \in [0, 2\pi)$ from $\hat{r}_1, \dots, \hat{r}_n$.
- 5: Order the set of angles $\hat{\theta}_1, \dots, \hat{\theta}_n$ in increasing order, and denote the induced ordering by $s = s_1, \dots, s_n$.
- 6: Compute the best circular permutation σ of the above ordering s that minimizes the resulting number of upsets with respect to the initial rank comparisons given by C

$$\sigma = \arg \min_{\sigma_1, \dots, \sigma_n} \|\text{sign}(P_{\sigma_i}(s)) - \text{sign}(C)\|_1 \quad (22)$$

with P defined as in (18).

- 7: Output as a final solution the ranking induced by the circular permutation σ .
-

is a comparison of the rankings obtained by the different methods: SVD, LS, SER, SER-GLM, RC, and SYNC-EIG, for an instance of the ranking problem given by the Erdős-Rényi

4. The circular or cyclic shift σ is given by $\sigma(i) = (i+1) \bmod n$. The result of repeatedly applying circular shifts to a given n -tuple $[x_1, \dots, x_n]$ are often denoted circular shifts of the tuple.

5. This is just the l_1 norm of the vectorized form of the matrix $\|X\|_1 = \sum_{i=1}^n \sum_{j=1}^n |X_{ij}|$

measurement graph $G(n = 100, p = 0.5)$ with cardinal comparisons, and outliers chosen uniformly at random with probability $\eta = \{0.35, 0.75\}$, according to the ERO(n, p, η) noise model (29). At low levels of noise, all methods yield satisfactory results, but as the noise level increases, only Sync-Rank is able to recover a somewhat accurate solution, and significantly outperforms all of the other methods in terms of the number of flips (i.e., the Kendall distance) with respect to the original ranking.

From a computational perspective, we point out that in most practical applications where the matrix of pairwise measurements is sparse, the eigenvector computation on which Sync-Rank is based can be solved in almost linear time. Every iteration of the power method is linear in the number of edges in the graph G (i.e., the number of pairwise measurements), but the number of iterations is greater than $O(1)$, as it depends on the spectral gap.

As a final remark, we point out that existing theoretical guarantees on the angular synchronization problem [4], [5], [30] imply lower bounds on the largest amount of noise permissible in the pairwise rank offsets while still achieving exact recovery of the ground truth ranking. This is due to the fact that perfect recovery of the angles in the angular synchronization problem is not necessary for a perfect recovery of the underlying true ranking, it only suffices that the relative ordering of the angles is preserved. However, the numerical results shown in the heat map (c) of Figure 4, corresponding to the outliers noise model on the pairwise cardinal rankings (as given by (29)), suggests that the threshold probability in the case of ranking is very similar to its analogue from angular synchronization obtained via the spectral relaxation, i.e., $O\left(\frac{1}{\sqrt{np}}\right)$ obtained by Singer [30] and illustrated by the overlaid black curve. In other words, as soon as angular synchronization fails, it does so by not only estimating incorrectly the magnitude of each and every angle, but also fails in preserving the relative ordering of the angles, which is precisely what Sync-Rank relies on.

4.1 Synchronization Ranking for Ordinal Comparisons

When the pairwise comparisons are ordinal, and thus $C_{ij} = \pm 1, \forall ij \in E$, all the angle offsets in the synchronization problem will have constant magnitude, which is perhaps undesirable. To this end, we report on a scoring method that associates a magnitude to each ± 1 entry, and gives a more accurate description of the rank-offset between a pair of players, which ultimately constitutes the input to Sync-Rank. For an ordered pair (i, j) , we define the *Superiority Score* of i with respect to j in a similar manner as the S_{ij}^{match} score (41) used by the Serial-Rank algorithm. Let

$$W_{ij} = \mathcal{L}(i) \cap \mathcal{H}(j) = \{k \mid C_{ik} = -1 \text{ and } C_{jk} = 1\} \quad (23)$$

where $\mathcal{L}(x)$ denotes the set of nodes with rank lower than x , and $\mathcal{H}(x)$ denotes the set of nodes with rank higher than x . The rank offset used as input for Sync-Rank is given by

$$S_{ij} = W_{ji} - W_{ij} \quad (24)$$

The philosophy behind this measure is as follows. In the end, we want S_{ij} to reflect the true rank-offset between two nodes. One can think of W_{ij} as the number of witnesses favorable to i (*supporters* of i), which are players ranked lower than i but higher than j . Similarly, W_{ji} is the number

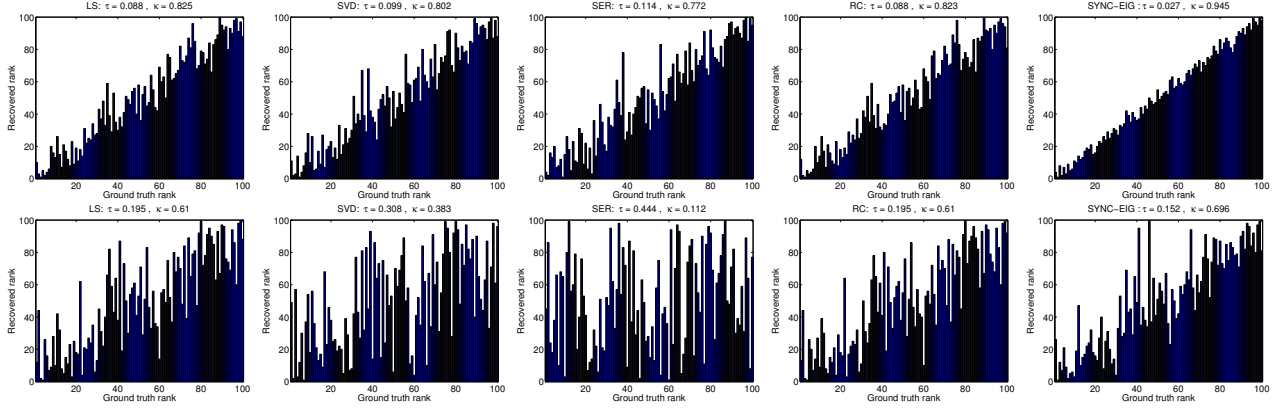


Fig. 2. Comparison of the methods for the ERO($n = 100, p = 0.5, \eta$) noise model (29) with cardinal comparisons with $\eta = 0.35$ (top), $\eta = 0.75$ (bottom).

of witnesses favorable to j (*supporters of j*), which are players ranked lower than j but higher than i . The final rank offset is given by the difference of the two values W_{ij} and W_{ji} (though one could also perhaps consider the maximum), and S_{ij} which can be interpreted as a proxy for $r_i - r_j$. We solve the resulting synchronization problem via the eigenvector method, and denote the approach by SYNC-SUP. While this approach yields rather poor results (compared to the initial synchronization method) on the synthetic data sets, its performance is comparable to that of the other methods when applied to the small sized Premier League soccer data set. However, it performs remarkably well on the NCAA basketball data set, and achieves a number of upsets that is twice as low as that of the second best method, as shown in Figure 8 of Section 5.4. One could investigate whether there is a connection between the rather intriguing performance of SYNC-SUP and the observation that for this particular data set, the measurement graphs for each season are somewhat close to having a one dimensional structure, with teams grouped in leagues according to their strengths, and it is less likely for a highly ranked team to play a match against a very weak team. It would be interesting to investigate the performance of all algorithms on a measurement graph that is a disc graph with nodes lying on a one-dimensional line corresponding to the ground truth ranks, with an edge between two nodes if and only if they are at most r units apart.

5 NOISE MODELS AND EXPERIMENTAL RESULTS

We compare the performance of Sync-Rank with that of other methods across a variety of measurement graphs, varying parameters such as the number of nodes, edge density, level of noise and noise models. We detail below the two noise models we consider, and then proceed with the performance outcomes of extensive numerical simulations.

5.1 Measurement and Noise Models

In most real world scenarios, noise is inevitable and imposes additional significant challenges, aside from those due to sparsity of the measurement graph. To each player i (corresponding to a node of G), we associate a corresponding unique positive integer weight r_i . For simplicity, one may choose to think of r_i as taking values in $\{1, 2, \dots, n\}$, and assume there is an underlying ground truth ranking of the

n players, with the most skillful player having rank 1, and the least skillful one having rank n . We denote by

$$\text{Cardinal measurement : } O_{ij} = r_i - r_j \quad (25)$$

the ground truth **cardinal** rank-offset of a pair of players. Thus, for cardinal comparisons, the available measurements C_{ij} are noisy versions of the ground truth O_{ij} entries. On the other hand, an **ordinal** measurement is a pairwise comparison between two players which reveals only who the higher-ranked player is, or in the case of items, the preference relationship between the two items, i.e., $C_{ij} = 1$ if item j is preferred to item i , and $C_{ij} = -1$ otherwise

$$\text{Ordinal measurement : } O_{ij} = \text{sign}(r_i - r_j). \quad (26)$$

without revealing the intensity of the preference relationship. This setup is commonly encountered in classical social choice theory, under the name of *Kemeny* model, where O_{ij} takes value 1 if player j is ranked higher than player i , and -1 otherwise. Given (perhaps noisy) versions of the pairwise cardinal or ordinal measurements O_{ij} given by (25) or (26), the goal is to recover an ordering (ranking) of all players that is as consistent as possible with the given data. We compare our proposed ranking methods with those summarized in Section 2, on measurement graphs of varying edge density, under two different noise models, and for both ordinal and cardinal comparisons.

To test for robustness against incomplete measurements, we use a measurement graph G of size n given by the popular Erdős-Rényi model $G(n, p)$, where edges between the players are present independently with probability p . We consider two different values of $p = \{0.2, 1\}$, the latter case corresponding to the complete graph K_n on n nodes, when comparisons between all possible pairs of players are available. To test the robustness against noise, we consider the following two noise models detailed below. We remark that, for both models, noise is added such that the resulting measurement matrix remains skew-symmetric.

5.1.1 Multiplicative Uniform Noise (MUN) Model

In the **Multiplicative Uniform Noise** model, which we denote by $\text{MUN}(n, p, \eta)$, noise is multiplicative and uniform, meaning that, for cardinal measurements, instead of the true rank-offset measurement $O_{ij} = r_i - r_j$, we actually measure

$$C_{ij} = O_{ij}(1 + \epsilon), \quad \text{where } \epsilon \sim [-\eta, \eta]. \quad (27)$$

Acronym	Name	Section
SVD	SVD Ranking	Sec. 2.3
LS	Least Squares Ranking	Sec. 2.4
SER	Serial-Ranking	Sec. 2.1
SER-GLM	Serial-Ranking in the GLM model	Sec. 2.1
RC	Rank-Centrality	Sec. 2.2
SYNC	Sync-Rank, spectral relaxation	Sec. 4
SYNC-SUP	Sync-Rank using the Sup. Score	Sec. 4.1
SYNC-SDP	Sync-Rank, SDP relaxation	Sec. 3.2

TABLE 1

Names of the algorithms, their acronyms, and respective Sections.

An equivalent formulation is that to each true rank-offset measurement we add random noise ϵ_{ij} (depending on O_{ij}) uniformly distributed in $[-\eta O_{ij}, \eta O_{ij}]$, i.e.,

$$C_{ij} = \begin{cases} r_i - r_j + \epsilon_{ij} & \text{for an existing edge} \quad \text{w. p. } p \\ 0 & \text{for a missing edge} \quad \text{w. p. } 1 - p. \end{cases} \quad (28)$$

with $\epsilon_{ij} \sim \mathcal{U}([-\eta(r_i - r_j), \eta(r_i - r_j)])$, where \mathcal{U} denotes the discrete uniform distribution. Note that we cap the erroneous measurements at $n - 1$ in absolute value. Thus, whenever $C_{ij} > n - 1$ we set it to $n - 1$, and whenever $C_{ij} < -(n - 1)$ we set it to $-(n - 1)$, since the furthest away two players can be is $n - 1$ positions. The percentage noise added is 100η , (e.g., $\eta = 0.1$ corresponds to 10% noise). Thus, if $\eta = 50\%$, and the ground truth rank offset $O_{ij} = r_i - r_j = 10$, then the available measurement C_{ij} is a random number in $[5, 15]$.

5.1.2 Erdős-Rényi Outliers (ERO) Model

The second noise model we consider is an **Erdős-Rényi Outliers** model, abbreviated by $\text{ERO}(n, p, \eta)$, where the available measurements are given by the following mixture

$$C_{ij} = \begin{cases} r_i - r_j & \text{w.p. } (1 - \eta)p \\ \sim \mathcal{U}[-(n - 1), n - 1] & \text{w.p. } \eta p \\ 0 & \text{w.p. } 1 - p, \end{cases} \quad (29)$$

Note that this noise model is similar to the one considered by Singer [30], in which angle offsets are either perfectly correct with probability $1 - \eta$, and randomly distributed on the unit circle, with probability η . We expect that in most practical applications, the first noise model (MUN) is perhaps more realistic, where most rank offsets are being perturbed by noise (to a smaller or larger extent, proportionally to their magnitude), as opposed to having a combination of perfectly correct rank-offsets and randomly chosen rank-offsets.

5.2 Numerical comparison on synthetic data sets

This section compares the spectral and SDP relaxations of Sync-Rank with the other methods summarized in Sec. 2, and also in Table 1. We measure the accuracy of the recovered solutions, using the popular Kendall distance, i.e., the fraction of (ordered) pairs of candidates that are ranked in different order (flips), in the two permutations, the original one and the recovered one. We compute the Kendall distance on a logarithmic scale (\log_2), and remark that in the error curves, for small noise levels, the error levels for certain methods are missing due to being equal to 0.

The heatmaps in Figure 4 compute the recovery error, as given by the Kendall distance between the ground truth and the rankings obtained via the spectral relaxation of Sync-Rank, as we vary the sparsity p in the measurement graph and the noise level in the data, for both cardinal and ordinal measurements, and under both the MUN (28) and ERO (29) noise models. For the ERO model, we overlay on the heatmap the curve given by the threshold probability for recovery $O\left(\frac{1}{\sqrt{np}}\right)$ by Singer [30], and note that in the cardinal case, it matches very well the empirical performance.

In Figure 5 we compare the methods for the case of cardinal pairwise measurements, while in Figure 6 we do so for ordinal measurements. We average all the results over 20 experiments. The SYNC and SYNC-SDP methods are far superior to the other methods in the case of cardinal measurements, and similar to the other methods for ordinal measurements. In Figures 3 we plot the recovered rank of the SDP program (17). Note that for favorable levels of noise, the SDP solution is indeed of rank 1 even if we did not specifically enforce this constraint, a phenomenon explained only recently in the work of Bandeira et al. [4], who investigated the tightness of this SDP relaxation.

5.3 Numerical Comparison on the English Premier League Data Set

The first real data set we consider is from sports data, in particular, several years of match outcomes from the **English Premier League Soccer**. We consider the three complete seasons 2011-2014, both home (C^{home}_{ij}) and away (C^{away}_{ij}) games. We pre-process the data and consider four methods to extract information from the game outcomes and build the pairwise comparison matrix C

- 1) *Total-Point-Difference* $C^{tpd}_{ij} = C^{home}_{ij} + C^{away}_{ij}$: for each pair of teams, we aggregate the total score of the two matches between teams i and j
- 2) *Sign-Total-Point-Difference* $C^{stpd}_{ij} = \text{sign}(C^{tpd}_{ij}) = \text{sign}(C^{home}_{ij} + C^{away}_{ij})$: considers the winner after aggregating the above score
- 3) *Net Wins* $C^{nw}_{ij} = \text{sign}(C^{home}_{ij}) + \text{sign}(C^{away}_{ij})$: the number of times a team has beaten the other one.
- 4) *Sign(Net Wins)* $C^{snw}_{ij} = \text{sign}(C^{nw}_{ij})$: only considers the winner in terms of the number of victories.

Note that C^{tpd} and C^{nw} lead to cardinal measurements, while C^{stpd} and C^{snw} to ordinal ones. One may interpret the above variations of the pre-processing step as a different winning criterion for the game under consideration, and at the same time, as a different opportunity to test and compare the algorithms. We remind the reader that for soccer games, a win is compensated with 3 points, a tie with 1 point, and a loss with 0 points, and the final ranking of the teams is determined by the cumulative points a team gathers throughout the season. Given a particular criterion, we are interested in finding an ordering that minimizes the number of upsets. We denote by \hat{r}_i the estimated rank of player i as computed by the method of choice. Recall that lower values of \hat{r}_i correspond to higher ranks (better players or preferred items). We then construct the induced (possibly incomplete) matrix of induced pairwise rank-offsets $\hat{C}_{ij} = \hat{r}_i - \hat{r}_j$, if $(i, j) \in E(G)$, and 0 otherwise, and remark that $\hat{C}_{ij} < 0$ denotes that the rank of player i is higher than

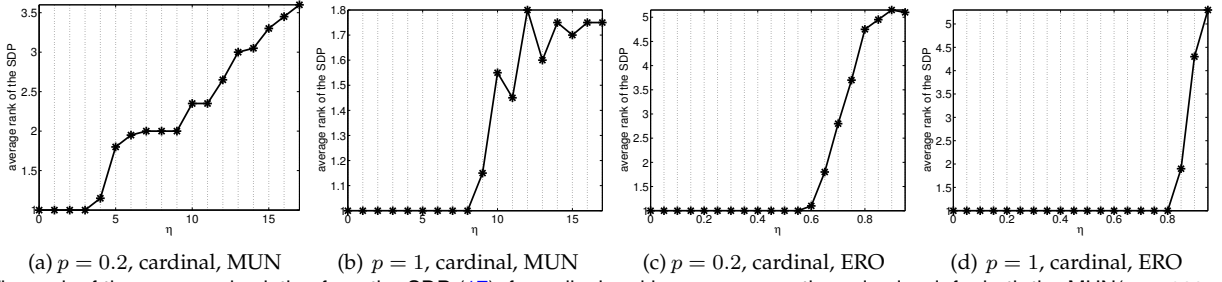


Fig. 3. The rank of the recovered solution from the SDP (17), for ordinal ranking, as we vary the noise level, for both the MUN($n = 200, p, \eta$) and ERO($n = 200, p, \eta$) noise models. We average the results over 20 runs. Note that if the SDP relaxation admits a rank-one solution, then the relaxation is tight and the synchronization problem was solved optimally.

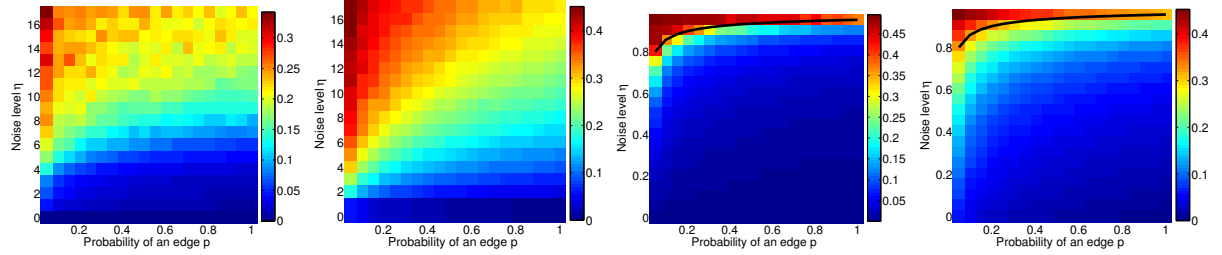


Fig. 4. The outcome of numerical experiments for the spectral relaxation of Sync-Rank, as we vary both the sparsity p in the measurement graph and the noise level η in the data, for both cardinal (left) and ordinal (right) measurements, and under both the MUN (top) and ERO (bottom) noise models, given by (28), respectively, (29). The color of the heatmaps encodes the Kendall distance between the estimated and the ground truth ranking. We fix $n = 500$ and average the results over 10 runs. For the ERO model, we overlay on the heatmap the curve given by the threshold probability for recovery $O(\frac{1}{\sqrt{np}})$ due to Singer [30], which, for cardinal measurements, is in excellent agreement with the empirical findings.

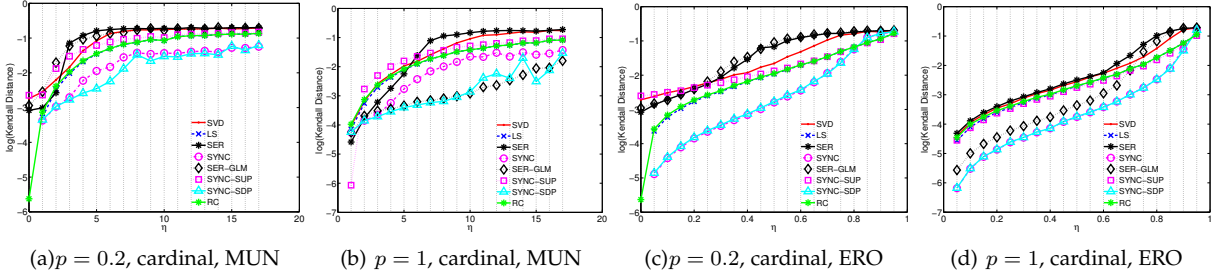


Fig. 5. Comparison of all ranking algorithms in the case of cardinal comparisons for the MUN($n = 200, p, \eta$) and ERO($n = 200, p, \eta$) noise models.

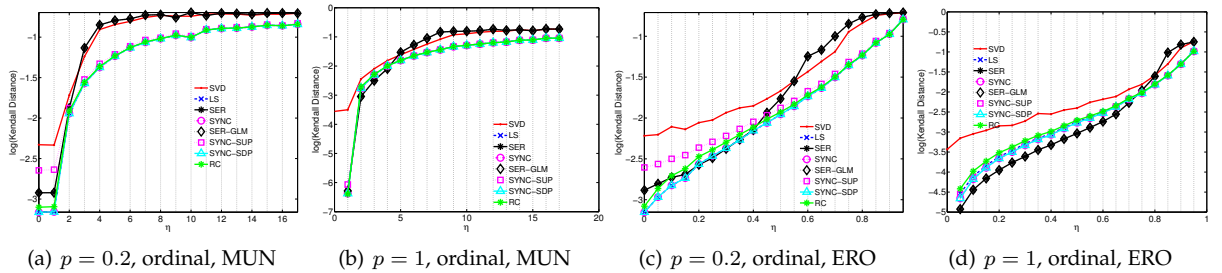


Fig. 6. Comparison of all methods for ranking with ordinal comparisons for the MUN($n = 200, p, \eta$) and ERO($n = 200, p, \eta$) noise models.

the rank of player j . To measure the accuracy of a proposed reconstruction, we rely on the popular metric

$$Q^{(u)} = \sum_{i=1}^{n-1} \sum_{j=i+1}^n \mathbf{1}_{\{\text{sign}(C_{ij} \hat{C}_{ij}) = -1\}} \quad (30)$$

which counts the number of disagreeing ordered comparisons. It contributes with a +1 to the summation whenever the ordering in the provided data contradicts the ordering in the induced ranking.

We show in Table 2 the rankings obtained by the different methods we have experimented with, for the 2013-2014 Premier League Season, when the input is based on the C^{mw} measurements (*Net Wins*). The final column, denoted by GT, shows the final official ranking at the end of the season. We sort the teams alphabetically, and show their assigned rank by each of the methods. The very last row in the Table computes the Kendall correlation between the respective ranking and the official final standing GT.

We remark that, across the different type of inputs considered, LS, SYNC and SYNC-SDP correlate best with the

official ranking. In terms of the number of upsets, the synchronization based methods alternate the first place in terms of quality with SER, depending on the pre-processing step and the input used for the ranking procedure. In addition, we show in Figure 7 the Q-scores associated to three recent seasons in Premier League: 2011-2012, 2012-2013, and 2013-2014, together with the mean across these three seasons. We show the number of upsets in Figure 7 and note that for

- C^{nw} : SYNC-SDP is the best performer, while the second place is tied by SYNC, SYNC-SUP and RC
- C^{snw} : LS, SYNC and SYNC-SDP are all tied for the first place
- C^{tpd} : SYNC-SUP is by far the best performer, followed by SER
- C^{stpd} : LS, SYNC and SYNC-SDP are all tied for the first place, followed by SER and SER-GLM on the second place.

Overall, we can conclude that, already for a very small example of size only $n = 20$, the synchronization-based methods show superior performance especially in the case of cardinal inputs C^{nw} and C^{tpd} .

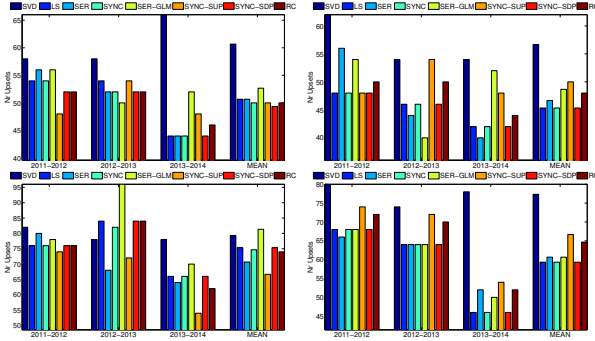
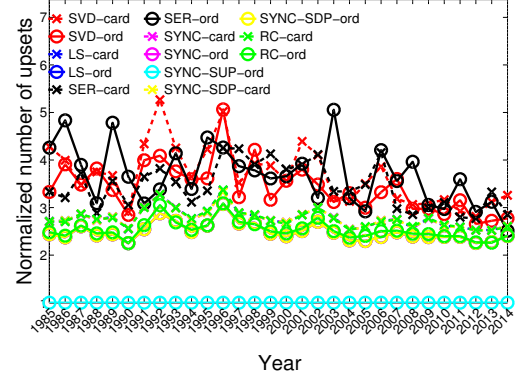


Fig. 7. Number of upsets $Q(u)$ (lower is better) for the English Premier League 2011-2014 data set, where the input is given by C^{nw} (top left) C^{snw} (top right), C^{tpd} (bottom left), and C^{stpd} (bottom right).

5.4 Numerical Comparison on the NCAA College Basketball Data Set

Our second real data set contains the outcomes of NCAA College Basketball matches during the regular season, for each of the past 30 seasons over the interval 1985 - 2014. During the regular season, most often it is the case that a pair of teams play against each other at most once. For example, during the 2014 season, for which 351 teams competed, there were a total of 5362 games, with each team playing on average about 30 games. We remark that in the earlier years, there were significantly fewer teams participating in the season, and therefore games played, which also explains the increasing number of upsets, given by (30), for the more recent years, as shown in Figure 8 (a). For example, during the 1985 season, only 282 teams participated, playing a total of 3737 games. In Figure 8 (a) we compare the performance of all methods across the years, for the case of both cardinal and ordinal measurements, i.e., when we consider the point difference or simply the winner of the game. Similarly, in Figure 8 (b) we compute the average number of upsets across all years in the interval 1985-2014, for both cardinal and ordinal measurements. Note that, for the less frequent

cases when two teams play against each other more than once, we simply average out the available scores. We remark that the SYNC-SUP method, i.e., eigenvector-based synchronization based on the superiority score from Section 4.1, performs remarkably well and achieves a number of upsets which is half that of the next best performing methods. The second best results are obtained (in no particular order) by LS, SYNC, SYNC-SDP and RC, which all yield very similar results, while SVD and SER are clearly the worst performers. We left out from the simulations the SER-GLM method due to its $O(n^3)$ computational running time.



(a) Number of upsets (normalized, for each year, by the value of the best performing method).

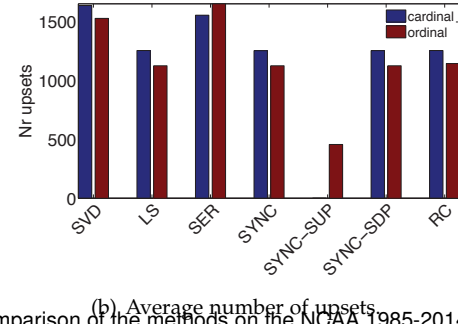


Fig. 8. Comparison of the methods on the NCAA 1985-2014 Basketball data set, in terms of the number of upsets (30), across the years (top) and averaged over the entire period (bottom) for both cardinal and ordinal measurements.

6 RANK AGGREGATION

In many practical instances, it is often the case that multiple voters or rating systems provide incomplete and inconsistent rankings or pairwise comparisons between the same set of players or items (e.g., at a contest, two players compete against each other, and k judges decide what the score should be). In such settings, a natural question to ask is, given a list of k (possibly incomplete) matrices of size $n \times n$ corresponding to ordinal or cardinal rank comparisons on a set of n players, how should one aggregate all the available data and produce a single global ranking, that is as consistent as possible with the observed data? The difficulty of the problem is amplified on one hand by the sparsity of the measurements, since each rating system only provides a comparison for a small set of pairs of items, and on the other hand by the amount of inconsistent information in the provided individual preferences. For example, given a set of 3 items A,B,C, the first judge may prefer A to B, the second one B to C, and the third one C to A.

Assume there are k judges, each of which makes available a noisy incomplete matrix $C^{(i)}$, $i = 1, \dots, k$ on the pairwise comparisons between a (perhaps randomly chosen) set of pairs of players. In practice, it is not necessarily the case that the same set of players appears in each comparison matrix $C^{(i)}$; however, for simplicity, we do assume this is the case. Furthermore, the pattern of pairwise existing and missing comparisons may be, and usually is, different across the above sequence of comparison matrices. We denote by $\Theta^{(i)}$, $i = 1, \dots, k$ the resulting matrices of angle offsets after the mapping (21), and by $H^{(i)}$, $i = 1, \dots, k$, the corresponding matrices after the mapping (9).

Tentative approach A. One possible naive approach to this problem would be to solve in parallel, via a method of choice, each of the $C^{(i)}$, $i = 1, \dots, k$ matrices, thus producing k possible rankings, one for each rating system, average the resulting rankings and consider the induced ordering by each method. In Figure 9 we denote by SVD-PAR, LS-PAR, SER-PAR, SYNC-PAR, SER-GLM-PAR, SYNC-SDP-PAR this approach for rank aggregation which runs, individually on each matrix $C^{(i)}$, $i = 1, \dots, k$ the various methods considered so far, and finally averages⁶ out the obtained rankings across all rankings proposed by a given method.

Tentative approach B. Another approach would be to first average out all the available matrices $C^{(i)}$, $i = 1, \dots, k$ into $\bar{C} = \frac{C^{(1)} + C^{(2)} + \dots + C^{(k)}}{k}$, and extract a final ranking by whatever preferred method that takes as input matrix \bar{C} , which we denote by $\bar{\Theta}$ after applying the transformation (21) In Figure 9 we denote by SVD-AVG, LS-AVG, SER-AVG, SYNC-AVG, SER-GLM-AVG, SYNC-SDP-AVG the approach for rank aggregation which runs the various methods on \bar{C} .

6.1 Rank Aggregation via Synchronization

A less naive approach to the above rank aggregation problem would be to consider the block diagonal matrix of size $N \times N$, with $N = nk$, given by $C_{N \times N} = \text{blkdiag}(C^{(1)}, C^{(2)}, \dots, C^{(k)})$, and its counterpart, the Hermitian matrix $H_{N \times N}$ obtained after mapping to the circle. Denote by \mathcal{A}_i , $i = 1, \dots, n$ the set of indices corresponding to the same player i in C . For simplicity, assume $\mathcal{A}_1 = \{1, n+1, 2n+1, \dots, (k-1)n+1\}$, and in general, $\mathcal{A}_i = \{i, n+i, 2n+i, \dots, (k-1)n+i\}$. As a first step, one could consider the SDP relaxation (17) which allows us to easily incorporate the hard constraints $\Upsilon_{ij} = 1$ if $i, j \in \mathcal{A}_u$, $\forall u = 1, \dots, k$, and add a post-processing step to ensure that the resulting top eigenvector v_1^Υ is piecewise constant along the sets of indices \mathcal{A}_u that represent the same players (since after the rank-1 projection of the solution Υ , the imposed hard constraints are no longer exactly satisfied). However, the size of the resulting SDP is prohibitive $nk \times nk$, and one may wonder if it is possible to reduce its size to $n \times n$, since there is a lot of redundant information in the formulation, with many subsets of the rows being equal, since they correspond to the same player. For computational reasons, it is also desirable to solve the rank aggregation problem via the eigenvector method. Indeed this is the

case, and this is based on the observation that the objective function for the rank-aggregation problem can be written as

$$\sum_{ij \in E(G)} e^{-\iota \theta_{ij}} \left(\sum_{u=1}^k e^{\iota \Theta_{ij}^{(u)}} \right) e^{\iota \theta_{ij}} = \sum_{ij \in E(G)} e^{-\iota \theta_{ij}} \bar{H}_{ij} e^{\iota \theta_{ij}} \quad (31)$$

where \bar{H} is given by the sum of the k Hermitian matrices $\bar{H} = \sum_{u=1}^k H^{(u)}$. The gist of this approach, leading to a much smaller-sized problem than the above SDP of size $nk \times nk$, is that we first average out (in the complex plane) all measurements corresponding to the same pair of players into a single $n \times n$ matrix \bar{H} , and then proceed with solving the ranking problem via the spectral and SDP relaxations of Sync-Rank, denoted by EIG-AGG, respectively SDP-AGG.

6.2 Numerical Results for Rank Aggregation

In this section, we provide numerical experiments comparing the various methods that solve the rank aggregation problem. Note that for ordinal measurements, we also compare to the original version of the Rank-Centrality algorithm, considered in [27] in the case of multiple rating systems and applicable only to ordinal measurements, whose main steps are summarized in (45), (46) and (47). We denote this approach by RCO in Figure 9. Furthermore, for both cardinal and ordinal measurements, we also compare to our proposed versions of Rank-Centrality, discussed in Sections C.1 and C.2 for the case of a single rating system $k = 1$. We adjust these two methods to the setting of multiple rating systems $k > 1$, by simply averaging out the winning probabilities in (50) or (51) of each rating system, as in (48).

In the top of Figure 9, we compare the methods in the setting of the Multiplicative Uniform Noise MUN($n = 100, p, \eta$), and average the results over 10 experiments. For cardinal data, we note the SDP-AGG method yields significantly more accurate results than the rest of the methods, followed by EIG-AGG, and the naive aggregation methods SDP-AVG and SDP-PAR. In the case of the complete graph, the Serial-Rank method SER-GLM comes in second best, together with SYNC-SDP-PAR. For ordinal data, the four methods EIG-AGG, SDP-AGG, SYNC-SDP-PAR and SYNC-SDPAVG yield very similar results, and are more accurate than the rest of the methods, especially for sparse graphs.

For the Erdős-Rényi Outliers model ERO($n = 100, p, \eta$) with cardinal measurements, illustrated in Figure 9 EIG-AGG, SDP-AGG are the best performers, followed very closely by SYNC-PAR and SYNC-SDP-PAR. For the case of the complete graph, the SER-GLM comes in next in terms of performance, while all the remaining methods show a rather poor relative performance. The relative performance of all methods is completely different for the case of ordinal measurements, under the ERO model, where SER-AVG and SER-GLM-AVG are the best performers, especially at lower levels of noise. Furthermore, the gap between these two methods and all the remaining ones increases as the measurement graphs become denser. For the case of the complete graph with ordinal comparisons with outliers, SER-AVG and SER-GLM-AVG produce results that are 2-3 orders of magnitude more accurate compared to all other methods, at lower levels of noise $\eta \leq 0.2$. As a final observation, we point out that, in the case of outliers given by the ERO model, our proposed version of Rank-Centrality, denoted by RC,

6. We simply consider the rank of each player across the different k rankings, and compute the average. The final averaged ranking is then recomputed from the induced average rankings

introduced in Section C.1 and given by (50), performs far better than the original RCO method in [27].

7 RANKING WITH HARD CONSTRAINTS VIA SEMIDEFINITE PROGRAMMING

Next, we consider a semi-supervised ranking problem, and illustrate how one can enforce certain types of constraints. Suppose that the user has readily available information on the true rank of a subset of players, and would like to obtain a global ranking solution which obeys the imposed ranking constraints. To this end, we propose a modified Sync-Rank algorithm based on the SDP relaxation of the ASP, followed by a post-processing step. Ranking with *vertex constraints*, i.e., information on the true rank of a small subset of players, is the ASP with constraints, similar to synchronization over \mathbb{Z}_2 with constraints considered in our previous work [13]. We refer to the subset of players whose rank is already known as *anchors* and denote this set by \mathcal{A} , while the set of non-anchor players (referred to as *free* players) shall be denoted by \mathcal{F} . In the ASP, *anchors* are nodes whose corresponding group element in $SO(2)$ is known a-priori. Given our measurement graph $G = (V, E)$ with node set V corresponding to a set of n group elements composed of h anchors $\mathcal{A} = \{a_1, \dots, a_h\}$ with $a_i \in SO(2)$, and l non-anchor elements $\mathcal{F} = \{f_1, \dots, f_l\}$ with $f_i \in SO(2)$, with $n = h + l$, and edge set E of size m corresponding to an incomplete set of m pairwise angle offsets of the type $f_i - f_j$ or $a_i - f_j$, the goal is to estimate $\hat{f}_1, \dots, \hat{f}_l \in SO(2)$.

Since the ground truth rank of anchor players is known, we compute the ground truth rank offset for all pairs of anchor nodes, which we then map to the circle via (21), and enforce the resulting offsets as hard constraints in the SDP relaxation (17), thus leading to the program in (33), where the maximization is taken over all matrices $\Upsilon \succeq 0$ with

$$\Upsilon_{ij} = \begin{cases} e^{\iota(f_i - f_j)} & \text{if } i, j \in \mathcal{F} \\ e^{\iota(f_i - a_j)} & \text{if } i \in \mathcal{F}, j \in \mathcal{A} \\ e^{\iota(a_i - a_j)} & \text{if } i, j \in \mathcal{A}. \end{cases} \quad (32)$$

with ones on the diagonal $\Upsilon_{ii} = 1, \forall i = 1, \dots, n$

$$\begin{aligned} & \underset{\Upsilon \in \mathbb{C}^{n \times n}}{\text{maximize}} && \text{Trace}(H\Upsilon) \\ & \text{subject to} && \Upsilon_{ii} = 1, i = 1, \dots, n \\ & && \Upsilon_{ij} = e^{\iota(a_i - a_j)}, \text{ if } i, j \in \mathcal{A} \\ & && \Upsilon \succeq 0. \end{aligned} \quad (33)$$

While in the noiseless case, the SDP in (33) may return a rank-1 one solution, for noisy data the rank could be larger than 1, and we consider the best rank-1 approximation. We use the top eigenvector $v_1^{(\Upsilon)}$ of Υ to recover the estimated angles and the induced rankings, but only up to a global rotation, i.e., a circular permutation. With this observation in mind, we propose a post-processing step, depicted in Figure 10, that ensures that the anchor players obey their prescribed rank. We keep the ranks of the anchor players fixed, and sequentially apply a cyclic permutation of the free players on the available rankings, searching for the optimal cyclic shift that minimizes the number of upsets (19). To measure the accuracy of each candidate ranking s (i.e., a permutation of size n), we first compute the associated pairwise rank offsets $Z(s) = s \otimes \mathbf{1} - \mathbf{1} \otimes s$. Next, we

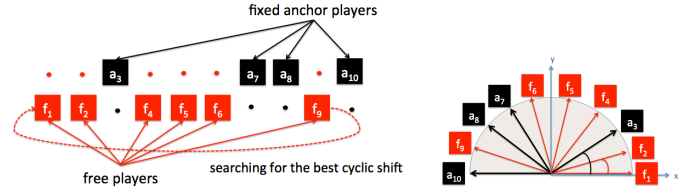


Fig. 10. Left: Searching for the best cyclic shift of the free players, shown in red. The black nodes denote the anchor players, whose rank is known and stays fixed. Right: The ground truth ranking.

choose the circular permutation $\sigma_{\mathcal{F}}^{(i)}$ of the ranks of the set of free players \mathcal{F} , which minimizes the l_1 norm of the following residual matrix that counts the total number of upsets $\frac{1}{2} \left\| \text{sign}(Z(\sigma_{\mathcal{A}} \sigma_{\mathcal{F}}^{(i)}) - \text{sign}(C)) \right\|_1$, where $\sigma_{\mathcal{A}}$ denotes the permutation associated to the anchor players, which stays fixed in the above minimization. The intuition is that the SDP solution preserves as best as possible the relative ordering of all players, in particular of the free players, and does so up to a global shift, which we recover via the above procedure. Figure 11 shows the outcome of nu-

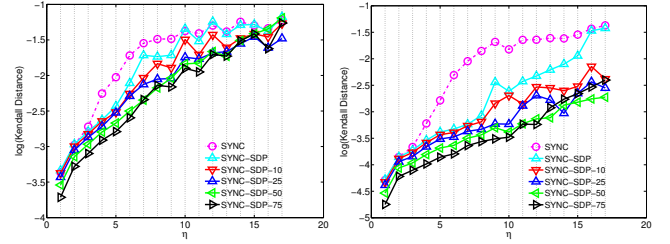


Fig. 11. Performance of Sync-Rank with constraints, for the MUN($n = 200, p, \eta$) noise model, with $p = 0.2$ (left) and $p = 1$ (right). SYNC-RANK denotes the SDP relaxation of Sync-Rank with K constraints.

merical simulations of Sync-Rank for several instances of the constrained ranking problem, under the MUN model with cardinal constraints. As expected, the numerical error decreases as more anchor information becomes available.

8 FUTURE RESEARCH DIRECTIONS

We highlight in this section several possible future research directions, which we expand in more detail in Appendix D. As seen in Section 7, adding pairwise rank-offsets as hard constraints improves the end results. A natural question to consider is how one may incorporate similar **soft** constraints, captured in a second matrix Q . As in (12), one can maximize the same quadratic form $z^* H z$, subject to the given constraints being satisfied well enough, i.e., $z^* Q z \geq \alpha$, and also $z^* z = n$, which altogether lead to a generalized eigenvalue problem. However, note that one could further replace H and Q with their respective graph Connection Laplacians, which are both PSD matrices, thus rendering the problem solvable in nearly-linear time due to recent progress in the area of Laplacian linear systems solvers, as explored in our recent work [11] on the well-studied constrained clustering problem. Another interesting variation concerns the scenario in which one wishes to enforce **ordinal hard constraints** that player i is ranked higher than player j , i.e., $\hat{r}_i > \hat{r}_j$. Another question is whether an

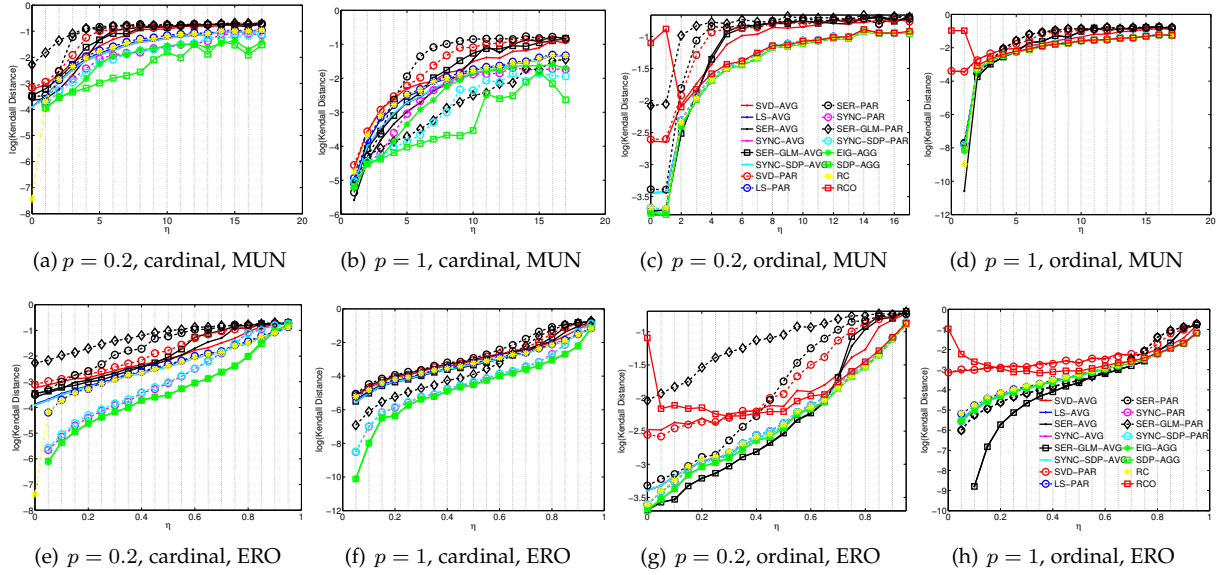


Fig. 9. Comparison of results for the Rank-Aggregation problem with $m = 5$ rating systems, for the Multiplicative Uniform Noise model MUN($n = 100, p, \eta$) (top) and the Erdős-Rényi Outliers model ERO($n = 100, p, \eta$) (bottom). We average the results over 10 runs.

l_1 -based formulation of the ASP could further improve the robustness to noise or sparsity of the measurement graph, since the existing spectral and SDP relaxations for the ASP [30] approximate a least squares solution in the angular space. One other very interesting direction concerns the possibility of casting the ranking problem as a synchronization problem over the group $SO(3)$, as opposed to $SO(2)$ as in present work. The extra *dimension* would facilitate the extraction of partial rankings, which one can collect by traversing along longitudinal lines and their immediate neighborhoods, and perhaps further increase the overall robustness to noise. Yet another future direction would be to aim for tighter guarantees in terms of robustness to noise. Finally, we have also considered the problem of planted partial rankings, which arises in the setting when a small subset Λ of players have their pairwise measurements more accurate than the rest of the network. We propose in Appendix A a spectral algorithm based on the Laplacian of the residual rank-offset matrix given by Sync-Rank, which, unlike other methods, is able to preserve the planted partial ranking and recover Λ , as shown in Figure 12.

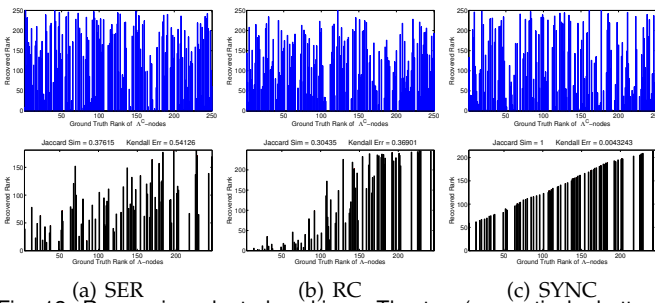


Fig. 12. Recovering planted rankings. The top (respectively, bottom) subplot corresponds to the recovered rankings of the non- Λ -players, respectively the Λ -players.

9 SUMMARY AND DISCUSSION

We considered the problem of ranking with noisy incomplete information and made an explicit connection with the

angular synchronization problem, for which spectral and SDP relaxations already exist in the literature with provable guarantees. This approach leads to a computationally efficient (as is the case for the spectral relaxation), non-iterative algorithm that is model independent and relies exclusively on the available data. We provided extensive numerical simulations, on both synthetic and real data sets that show that our proposed procedure compares favorably with state-of-the-art methods from the literature, across a variety of measurement graphs and noise models.

We took advantage of the spectral, and especially SDP, relaxations and proposed methods for ranking in the semi-supervised setting where a subset of the players have a prescribed rank to be enforced as a hard constraint. In addition, we considered the popular rank aggregation problem of combining ranking information from multiple rating systems that provide independent incomplete and inconsistent pairwise measurements on the same set of items, with the goal of producing a single global ranking.

REFERENCES

- [1] N. ALON, M. KRIVELEVICH, AND B. SUDAKOV, *Finding a large hidden clique in a random graph.*, in SODA, H. J. Karloff, ed., ACM/SIAM, 1998, pp. 594–598.
- [2] J. E. ATKINS, E. G. BOMAN, AND B. HENDRICKSON, *A spectral algorithm for seriation and the consecutive ones problem*, SIAM Journal on Computing, 28 (1998), pp. 297–310.
- [3] H. AZARI SOUFIANI, W. CHEN, D. C. PARKES, AND L. XIA, *Generalized Method-of-Moments for Rank Aggregation*, in Proceedings of the Annual Conference on Neural Information Processing Systems (NIPS 2013), 2013.
- [4] A. BANDEIRA, N. BOUMAL, AND A. SINGER, *Tightness of the maximum likelihood semidefinite relaxation for angular synchronization*, arXiv preprint arXiv:1411.3272, (2014).
- [5] A. S. BANDEIRA, A. SINGER, AND D. A. SPIELMAN, *A Cheeger inequality for the graph Connection Laplacian*, SIAM Journal on Matrix Analysis and Applications, 34 (2013), pp. 1611–1630.
- [6] F. BENAYCH-GEORGES AND R. R. NADAKUDITI, *The singular values and vectors of low rank perturbations of large rectangular random matrices*, Journal of Multivariate Analysis, 111 (2012), pp. 120–135.
- [7] N. BOUMAL, *A Riemannian low-rank method for optimization over semidefinite matrices with block-diagonal constraints*, arXiv preprint arXiv:1506.00575, (2015).

- [8] S. BOYD, N. PARIKH, E. CHU, B. PELEATO, AND J. ECKSTEIN, *Distributed optimization and statistical learning via the alternating direction method of multipliers*, Found. Trends Mach. Learn., 3 (2011), pp. 1–122.
- [9] R. A. BRADLEY AND M. E. TERRY, *Rank analysis of incomplete block designs: I. the method of paired comparisons*, Biometrika, (1952), p. 324345.
- [10] P. CREMONESI, Y. KOREN, AND R. TURRIN, *Performance of recommender algorithms on top-n recommendation tasks*, in Proceedings of the fourth ACM conference on Recommender systems, ACM, 2010, pp. 39–46.
- [11] M. CUCURINGU, I. KOUTIS, AND S. CHAWLA, *Scalable Constrained Clustering: A Generalized Spectral Method*, submitted, (2015).
- [12] M. CUCURINGU, Y. LIPMAN, AND A. SINGER, *Sensor network localization by eigenvector synchronization over the Euclidean group*, ACM Trans. Sen. Netw., 8 (2012), pp. 19:1–19:42.
- [13] M. CUCURINGU, A. SINGER, AND D. COWBURN, *Eigenvector synchronization, graph rigidity and the molecule problem*, Information and Inference, 1 (2012), pp. 21–67.
- [14] C. DE KERCHOVE AND P. VAN DOOREN, *Iterative filtering in reputation systems*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 1812–1834.
- [15] U. FEIGE AND J. R. LEE, *An improved approximation ratio for the minimum linear arrangement problem*, Inf. Process. Lett., 101 (2007), pp. 26–29.
- [16] U. FEIGE AND M. SELTZER, *On the densest k-subgraph problem*, Technical report CS97-16, Tel-Aviv University, Dept. of Computer Science, (1997).
- [17] D. FÉRAL AND S. PÉCHÉ, *The largest eigenvalue of rank one deformation of large Wigner matrices*, Communications in Mathematical Physics, 272 (2007), pp. 185–228.
- [18] F. FOGEL, A. D’ASPREMONT, AND M. VOJNOVIC, *Serialrank: Spectral ranking using seriation*, in Advances in Neural Information Processing Systems 27, 2014, pp. 900–908.
- [19] F. FOGEL, R. JENATTON, F. BACH, AND A. D’ASPREMONT, *Convex relaxations for permutation problems.*, in NIPS, C. J. C. Burges, L. Bottou, Z. Ghahramani, and K. Q. Weinberger, eds., 2013, pp. 1016–1024.
- [20] D. GLEICH, *SVD based term suggestion and ranking system*, in In: IEEE International Conference on Data Mining, 2004, pp. 391–394.
- [21] A. N. HIRANI, K. KALYANARAMAN, AND S. WATTS, *Least squares ranking on graphs*, arXiv preprint arXiv:1011.1716, (2010).
- [22] X. JIANG, L.-H. LIM, Y. YAO, AND Y. YE, *Statistical ranking and combinatorial hodge theory*, Mathematical Programming, 127 (2011), pp. 203–244.
- [23] D. R. KARGER, S. OH, AND D. SHAH, *Iterative Learning for Reliable Crowdsourcing Systems*, in Advances in Neural Information Processing Systems 24, J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, eds., Curran Associates, Inc., 2011, pp. 1953–1961.
- [24] R. LAI AND S. OSHER, *A splitting method for orthogonality constrained problems*, J. Sci. Comput., 58 (2014), pp. 431–449.
- [25] J. R. LEE, S. OVEIS GHARAN, AND L. TREVISAN, *Multi-way spectral partitioning and higher-order cheeger inequalities*, in Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing, STOC ’12, New York, NY, USA, 2012, ACM, pp. 1117–1130.
- [26] R. D. LUCE, *Individual Choice Behavior a Theoretical Analysis*, John Wiley and Sons, 1959.
- [27] S. NEGAHBAN, S. OH, AND D. SHAH, *Iterative ranking from pairwise comparisons*, in Advances in Neural Information Processing Systems 25, 2012, pp. 2474–2482.
- [28] B. OSTING, J. DARBON, AND S. OSHER, *Statistical ranking using the l1-norm on graphs*, vol. 7, 2013, pp. 907–926.
- [29] L. PAGE, S. BRIN, R. MOTWANI, AND T. WINOGRAD, *The pagerank citation ranking: Bringing order to the web*, in Proceedings of the 7th International World Wide Web Conference, 1998, pp. 161–172.
- [30] A. SINGER, *Angular synchronization by eigenvectors and semidefinite programming*, Appl. Comput. Harmon. Anal., 30 (2011), pp. 20–36.
- [31] A. SINGER AND H.-T. WU, *Vector diffusion maps and the connection laplacian*, Communications on Pure and Applied Mathematics, 65 (2012), pp. 1067–1144.
- [32] L. VANDENBERGHE AND S. BOYD, *Semidefinite programming*, SIAM Review, 38 (1994), pp. 49–95.
- [33] L. WANG AND A. SINGER, *Exact and stable recovery of rotations for robust synchronization*, Information and Inference, 2 (2013), pp. 145–193.
- [34] X. WANG, B. QIAN, AND I. DAVIDSON, *On constrained spectral clustering and its applications*, Data Min. Knowl. Discov., 28 (2014), pp. 1–30.
- [35] F. L. WAUTHIER, M. I. JORDAN, AND N. JOJIC, *Efficient ranking from pairwise comparisons.*, in ICML (3), vol. 28 of JMLR Proceedings, JMLR.org, 2013, pp. 109–117.
- [36] Z. WEN, D. GOLDFARB, AND K. SCHEINBERG, *Block coordinate descent methods for semidefinite programming*, in Handbook on Semidefinite, Conic and Polynomial Optimization, Springer, 2012, pp. 533–564.
- [37] L. XIONG AND L. LIU, *A reputation-based trust model for peer-to-peer e-commerce communities*, in E-Commerce, 2003. CEC 2003. IEEE International Conference on, IEEE, 2003, pp. 275–284.
- [38] S. YU, *Angular embedding: A robust quadratic criterion.*, IEEE Trans. Pattern Anal. Mach. Intell., 34 (2012), pp. 158–173.
- [39] S. ZHANG AND Y. HUANG, *Complex quadratic optimization and semidefinite programming*, SIAM Journal on Optimization, 16 (2006), pp. 871–890.

ACKNOWLEDGMENTS

The author would like to thank the Simons Institute for their warm hospitality and fellowship support during his stay in Berkeley throughout Fall 2014, when most of this research was carried out. He is very grateful to Amit Singer for suggesting the possibility of applying angular synchronization to the ranking problem. He would like to thank Andrea Bertozzi for her support via AFOSR MURI grant FA9550-10-1-0569, Alex d’Aspremont and Fajwel Fogel for useful discussions on Serial-Rank and relevant data sets, and Xiuyuan Cheng for discussions on random matrix theory. Finally, he would also like to thank the anonymous reviewers for their thorough comments and suggestions.



Mihai Cucuringu received his PhD degree at the Program in Applied and Computational Mathematics (PACM), Princeton University in 2012, and his BA degree in mathematics, computer science and economics from Hiram College, in 2007. He currently holds a postdoctoral position as a CAM Assistant Adjunct Professor in Computational Applied Mathematics, in the Department of Mathematics at University of California, Los Angeles. His current research focuses on the development and mathematical analysis of algorithms for large networks and data analysis, with applications to various problems in engineering, machine learning, and biology.

APPENDIX A

RECOVERING PLANTED PARTIAL RANKINGS

Aside from its increased robustness to noise, we point out another appealing feature of the Sync-Rank algorithm. In many real-life scenarios, one is not necessarily interested in recovering the entire ranking, but only the ranking of a subset of the players. Often, the noise in the available measurements may not be distributed uniformly throughout the network, with part of the network containing pairwise comparisons that are a lot less noisy than those in the rest of the network. This asymmetry provides us with an opportunity and raises the question of whether one can recover partial rankings which are locally consistent with the given data. In this Appendix we demonstrate empirically that Sync-Rank is able to preserve such partial rankings. In addition, we extract such hidden locally consistent partial rankings, relying on a spectral algorithm in the spirit of existing algorithms for the planted clique and densest subgraph problems.

We remark that this approach can be extended to extracting multiple partial rankings by relying on recent work of Lee et al. on multi-way spectral partitioning and higher-order Cheeger inequalities [25]. Note that in the case of k partial rankings, one does not necessarily seek a partition of the residual graph into k clusters, but rather seeks k disjoint clusters, whose union is not necessarily the entire graph. The *Cheeger Sweep* final step of their algorithm allows one to extract k such sets of smallest expansion. For simplicity, we only consider the case of a single planted locally consistent partial ranking. We denote by Λ the set of players whose pairwise ranking measurements contain less noise than the rest of the available data, and we shall refer to such players as Λ -players, and to the remaining ones as Λ^C -players.

We evaluate our proposed procedure on graphs generated from the following ensemble $\mathcal{G}(n, p, \beta, \eta_1, \eta_2)$, where the measurement graph G is an Erdős-Rényi random graph $G(n, p)$, $|\Lambda| = \beta n$, and the noise model is given by the following mixture

$$C_{ij} = \begin{cases} r_i - r_j & \{i, j\} \in \Lambda, & \text{w.p. } (1 - \eta_1)\alpha \\ \sim \mathcal{U}[-(n-1), n-1] & \{i, j\} \in \Lambda & \text{w.p. } \eta_1\alpha \\ r_i - r_j & \{i, j\} \notin \Lambda, & \text{w.p. } (1 - \eta_2)\alpha \\ \sim \mathcal{U}[-(n-1), n-1] & \{i, j\} \notin \Lambda & \text{w.p. } \eta_2\alpha \\ 0 & & \text{w.p. } 1 - \alpha. \end{cases} \quad (34)$$

where \mathcal{U} denotes the discrete uniform distribution.

The approach we propose for recovering such a locally consistent partial ranking starts as follows. For a given ranking method of choice (such as the ones we have considered throughout the paper), we estimate the complete ranking $\hat{r}_1, \dots, \hat{r}_n$, and consider the resulting matrix of rank offsets

$$\hat{C} = (\hat{r} \otimes \mathbf{1} - \mathbf{1} \otimes \hat{r}) \circ A \quad (35)$$

where \otimes denotes the outer product of two vectors $x \otimes y = xy^T$, \circ denotes the Hadamard product of two matrices (entrywise product), and A is the adjacency matrix of the graph G ; in other words $\hat{C}_{ij} = \hat{r}_i - \hat{r}_j$, $(i, j) \in E(G)$. Next, we consider the *residual matrix* of pairwise rank offsets

$$R = |C - \hat{C}| \quad (36)$$

for the case of cardinal measurements, and

$$R = |C - \text{sign}(\hat{C})| \quad (37)$$

for the case of ordinal measurements. We illustrate in Figure 13 the residual matrices obtained by each of the methods, and remark that, for ease of visualization, the set Λ consists of the first $\beta n = 75$ nodes, corresponding to the top left corner of each of the residual matrices shown in the heatmap. Whenever an estimated rank-offset \hat{C}_{ij} (induced by the recovered ranking) matches very well with the initial measurement C_{ij} , we expect the residual R_{ij} to have small magnitude, and conversely, whenever the recovered offset \hat{C}_{ij} is far from the initial measurement C_{ij} then we expect R_{ij} to have large magnitude. Furthermore, we show in Figure 14, the recovered rankings by each of the methods, where we separate the rankings of the Λ^C -players (shown in the top subplot, in blue), from those of the Λ -players (shown in the bottom subplot, in red). Note that the set Λ is not available to the user, and at this point we do not have yet an estimate $\hat{\Lambda}$ for Λ ; however this plot already highlights the fact that synchronization-based ranking preserves almost perfectly the relative ranking of the Λ -players, while all the other methods fail to do so. Note that the title of each of the bottom plots of Figure 14 shows the Jaccard Similarity index between Λ and $\hat{\Lambda}$ (which we show next how to estimate), and the Kendall distance between the ground truth rankings of the Λ -players, and the estimated rankings of the Λ -players.

In practice, when the measurements between pairs of Λ -nodes also contain noise, the corresponding sub-matrix has no longer zero or close to zero residuals, and it is harder to identify such a set of nodes. The resulting task at hand is to identify a subset of nodes for which the average inter-edge weights (the residuals shown in Figure 13) is as small as possible. This problem is equivalent⁷ to the well-known *densest subgraph* problem, investigated in the theoretical computer science literature, for the case when the graph G is unweighted. *Densest-k-Subgraph* (DkS) on an undirected unweighted graph G concerns finding a subset of nodes $U \in V(G)$ of size $|U| = k$ with the maximum induced average degree. Using a reduction from the Max-Clique problem, the *Densest-k-Subgraph* can be shown to be NP-hard to solve exactly. Feige and Seltzer [16] showed that the *Densest-k-Subgraph* problem is NP-complete even when restricted to bipartite graphs of maximum degree 3. If the parameter k is not known a-priori, the problem becomes the well known *Densest Subgraph* problem, where one seeks to find a subset U (regardless of its size) in order to maximize the average degree induced by U . The *Densest Subgraph* problem can be solved in polynomial time using either linear programming or flow techniques.

In a seminal paper [1], Alon, Krivelevich, and Sudakov considered the graph ensemble $G(n, 1/2, k)$, where one starts with an Erdős-Rényi $G(n, p = 1/2)$ and randomly places a clique of size k in G . They proposed an efficient spectral algorithm that relies on the second eigenvector of the adjacency matrix of G , which almost surely finds the clique of size k , as long as $k > c\sqrt{n}$. We remark that

⁷. Perhaps after hard thresholding the entries of the residual matrix to 0 or 1.

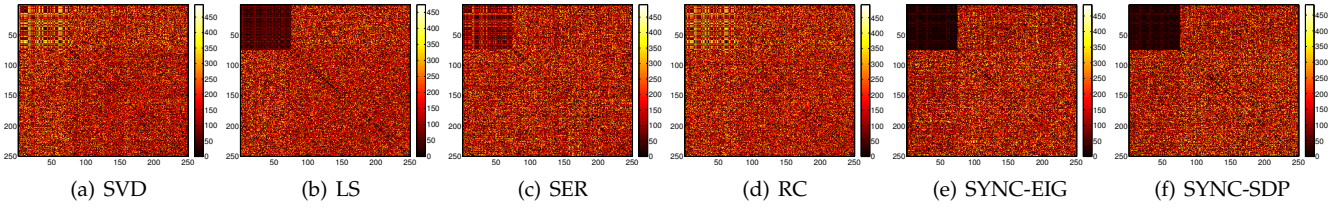


Fig. 13. The residual matrices given by (36), for each of the six methods, for a random instance of the ensemble $\mathcal{G}(n = 250, \beta = 0.3, \eta_1 = 0, \eta_2 = 1)$. The Λ -nodes correspond to the first $\beta n = 75$ positions.

in our case the residual graph is weighted and thus a different procedure is required, though one can perhaps think of various heuristics for thresholding the entries of W , setting to 1 all remaining weights (thus rendering the graph unweighted), and perhaps further adding edges to the graph between r -hop neighbors away pairs of nodes, with the hope of producing a clique, which we can then detect using the algorithm proposed in [1].

However, in this work, we take none of the above approaches, and rely on another spectral algorithm to uncover the planted dense subgraph, and hence the partial ranking. We use the following approach to uncover the planted densest subgraph of given size k . We first compute a similarity between the adjacent nodes of the graph using the Gaussian Kernel

$$W_{ij} = e^{-\frac{R_{ij}^2}{\epsilon^2}}, (i, j) \in E(G)$$

for the parameter choice $\epsilon^2 = 2n/10$, and consider the random-walk Laplacian $L = D^{-1}W$, whose trivial eigenvalue-eigenvector pair is given by $\lambda_1 = 1$ and the all-ones vector $v_1 = \mathbf{1}$. To extract a cluster of given size k (our proposed estimate for the set Λ), we compute the first non-trivial eigenvector v_2 of L , and consider its top k largest entries to produce an estimate $\hat{\Lambda}$ for the set Λ . We detail these steps in Algorithm 2. Figure 15 shows a comparison of the methods (results averaged across 15 experiments) in terms of the Jaccard Similarity Index between the ground truth set Λ and our proposed estimate $\hat{\Lambda}$, defined as

$$\mathcal{J}(\Lambda, \hat{\Lambda}) = \frac{|\Lambda \cap \hat{\Lambda}|}{|\Lambda \cup \hat{\Lambda}|}, \quad (39)$$

meaning that for values equal to 1 the algorithm would perfectly recover the planted nodes in Λ . Furthermore, we plot in Figure 16 the corresponding Kendall distance between the recovered ranking of the $\hat{\Lambda}$ -nodes (from Figure 15) and their ground truth values, for the ensemble given by $\mathcal{G}(n = 250, p, \beta, \eta_1, \eta_2 = 1)$, for varying $p = \{0.2, 0.5, 1\}$ and $\eta_1 = \{0, 0.2, 0.4\}$. Note that across all experiments, the SYNC-SDP algorithm is by far the best performer, both in terms of the Jaccard Index (thus being able to recover the set Λ) and the Kendall distance (thus being able to recover accurately the ranking⁸ of the estimated $\hat{\Lambda}$ -nodes). Note that an alternative approach to estimating the ranking of the $\hat{\Lambda}$ -nodes (once $\hat{\Lambda}$ has already been estimated) would be to recompute it by re-running the method of choice on the original measurement matrix C , restricted to the sub-matrix

8. This partial ranking is simply extracted from the complete ranking solution, as computed using the method of choice.

Algorithm 2 Algorithm for recovering the densest (weighted) subgraph of size k of the residual graph whose adjacency matrix is denoted by W .

Require: A weighted graph $H = (V, E)$, and its symmetric adjacency matrix W , and $k < n$

- 1: Compute the random-walk Laplacian $L = D^{-1}W$, where D is a diagonal matrix with $D_{ii} = \sum_{j=1}^n W_{ij}$
- 2: Find the second eigenvector v_2 of L , corresponding to the second largest eigenvalue $\lambda_2 < \lambda_1 = 1$
- 3: Sort the entries of v_2 in decreasing order, consider the top k largest entries in v_2 , let Λ^d denote the set of corresponding vertices, and compute the residual error associated to Λ^d

$$ERR^{\Lambda^d} = \|C_{\Lambda^d, \Lambda^d} - \hat{C}_{\Lambda^d, \Lambda^d}\|_1$$

- 4: Sort the entries of v_2 in increasing order, consider the top k largest entries in v_2 , let Λ^i denote the set of corresponding vertices, and compute the residual error associated to Λ^i

$$ERR^{\Lambda^i} = \|C_{\Lambda^i, \Lambda^i} - \hat{C}_{\Lambda^i, \Lambda^i}\|_1$$

- 5: Output the final estimate $\hat{\Lambda}$ for Λ as

$$\hat{\Lambda} = \arg \min_{\Lambda^d, \Lambda^i} \{ERR^{\Lambda^d}, ERR^{\Lambda^i}\} \quad (38)$$

corresponding to nodes $\hat{\Lambda}$. We expect this approach to yield more accurate results, especially in the case when the measurements between a pair of (Λ^C, Λ^C) -nodes, or a pair of (Λ, Λ^C) -nodes are very noisy (i.e., η_2 is large in the ensemble $\mathcal{G}(n, p, \beta, \eta_1, \eta_2)$ defined by (34)), as was the case for the numerical experiments detailed here, for which we have chosen $\eta_2 = 1$, meaning such measurements are pure noise and chosen uniformly at random from $[-(n-1), (n-1)]$.

APPENDIX B SERIAL-RANK

A spectral algorithm that exactly solves the noiseless seriation problem (and the related *continuous ones problem*) was proposed by Atkins et al. [2], based on the observation that given similarity matrices computed from serial variables, the ordering induced by the second eigenvector of the associated Laplacian matrix (i.e., the *Fiedler* vector) matches that of the variables. In other words, this approach (reminiscent

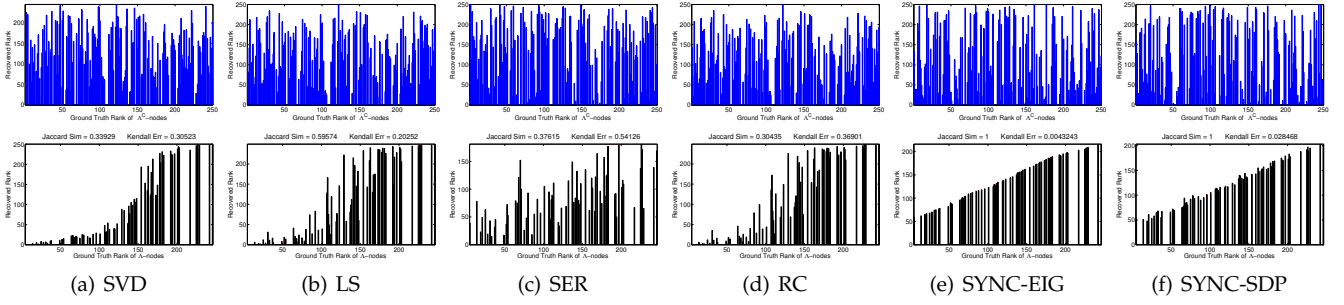


Fig. 14. The recovered rankings by each of the six methods, for a random instance of the ensemble $\mathcal{G}(n = 250, \beta = 0.3, \eta_1 = 0, \eta_2 = 1)$, for which $|\Lambda| = \beta n = 75$. For each plot, the top subplot, respectively bottom subplot, corresponds to the recovered rankings of the (ground truth) Λ^C -nodes, respectively the (ground truth) Λ -nodes. Note that, in practice, the set Λ is not known a-priori, but we use this information here only for the purpose of making the point that the synchronization-based method is able to perfectly preserve the relative ranking of the Λ -players. We exploit this phenomenon, and are able to recover the planted set Λ in a totally unsupervised manner.

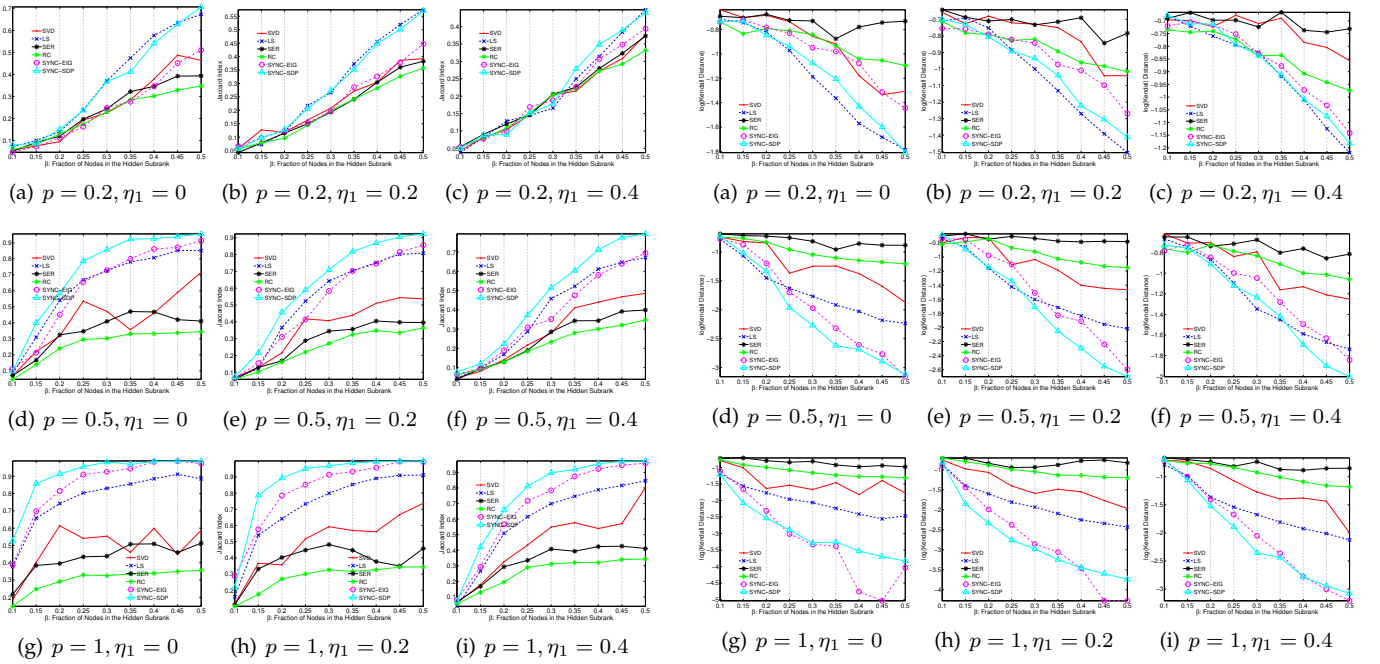


Fig. 15. Comparison of the methods of the Jaccard Similarity Index (higher is better) between the recovered $\hat{\Lambda}$ and the ground truth Λ , from the ensembles given by $\mathcal{G}(n = 250, p, \beta, \eta_1, \eta_2 = 1)$, for varying $p = \{0.2, 0.5, 1\}$ and $\eta_1 = \{0, 0.2, 0.4\}$. Experiments are averaged over 15 runs.

of spectral clustering) exactly reconstructs the correct ordering if the items in question lie on a one dimensional chain. In [18], Fogel et al. adapt the above seriation procedure to the ranking problem, and propose an efficient polynomial-time algorithm with provable recovery and robustness guarantees. Under certain conditions on the pattern of noisy entries, Serial-Rank is able to perfectly recover the underlying true ranking, even when a fraction of the comparisons are either corrupted by noise or completely missing. In this noisy setting, when the underlying measurement graph is dense, in other words, a high fraction of all the pairwise comparisons are observed, the spectral solution obtained by Serial-Rank is more robust to noise than other classical scoring-based methods.

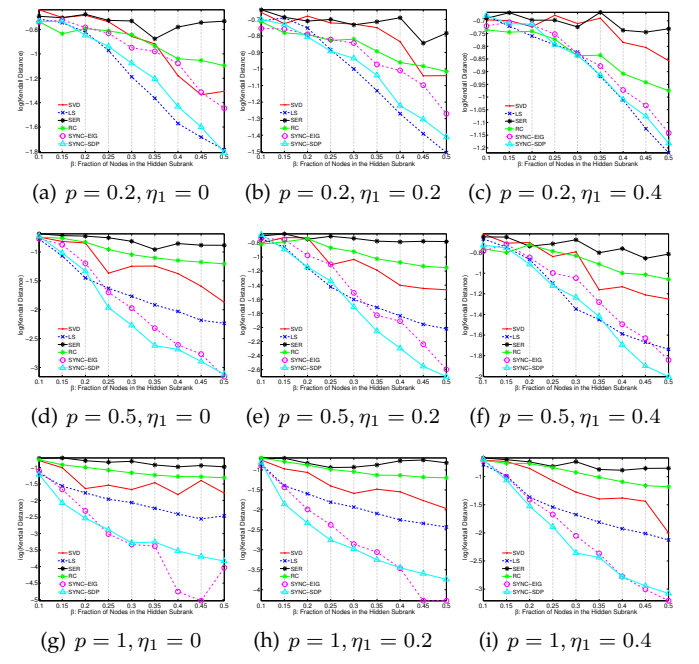


Fig. 16. The Kendall Distance (lower is better) between the recovered ranking of the $\hat{\Lambda}$ -nodes by each of the methods (from Figure 15) and their ground truth values, for the ensemble given by $\mathcal{G}(n = 250, p, \beta, \eta_1, \eta_2 = 1)$, for varying $p = \{0.2, 0.5, 1\}$ and $\eta_1 = \{0, 0.2, 0.4\}$. Experiments are averaged over 15 runs.

The authors of [18] propose two approaches for computing a pairwise similarity matrix from both ordinal and cardinal comparisons between the players. In the case of ordinal measurements, the similarity measure counts the number of *matching comparisons*. More precisely, given as input a skew symmetric matrix C of size $n \times n$ of pairwise comparisons $C_{ij} = \{-1, 0, 1\}$, with $C_{ij} = -C_{ji}$, given by the following model

$$C_{ij} = \begin{cases} 1 & \text{if } i \text{ is ranked higher than } j \\ 0 & \text{if } i \text{ and } j \text{ are tied, or never competed} \\ -1 & \text{if } j \text{ is ranked higher than } i \end{cases} \quad (40)$$

For convenience, the diagonal of C is set to $C_{ii} = 1, \forall i =$

1, 2, ..., n . Finally, the pairwise similarity matrix is given by

$$S_{ij}^{match} = \sum_{k=1}^n \left(\frac{1 + C_{ik}C_{jk}}{2} \right). \quad (41)$$

Note that $C_{ik}C_{jk} = 1$ whenever i and j have the same signs, and $C_{ik}C_{jk} = -1$ whenever they have opposite signs, thus S_{ij}^{match} counts the number of matching comparisons between i and j with a third reference item k . The lack of a comparison of item k with either i or j contributes with a $\frac{1}{2}$ to the summation in (41). The intuition behind this similarity measure proposed by Fogel et al. is that players that beat the same players and are beaten by the same players should have a similar ranking in the final solution. Written in a compact form, the final similarity matrix is given by

$$S^{match} = \frac{1}{2} (n\mathbf{1}\mathbf{1}^T + CC^T) \quad (42)$$

We summarize in Algorithm 3 the main steps of the Serial-Rank method of Fogel et al. for ranking via seriation.

Algorithm 3 Serial-Rank: an algorithm for spectral ranking using seriation [18].

Require: A set of pairwise comparisons $C_{ij} \in \{-1, 0, 1\}$ or $[-1, 1]$

- 1: Compute a similarity matrix as shown in (41)
- 2: Compute the associated graph Laplacian matrix

$$L_S = D - S \quad (43)$$

where D is a diagonal matrix $D = \text{diag}(S\mathbf{1})$, i.e., $D_{ii} = \sum_{j=1}^n G_{i,j}$ is the degree of node i in the measurement graph G .

- 3: Compute the Fiedler vector of S (eigenvector corresponding to the smallest nonzero eigenvalue of L_S).
 - 4: Output the ranking induced by sorting the Fiedler vector of S , with the global ordering (increasing or decreasing order) chosen such that the number of upsets is minimized.
-

In the generalized linear model setting, one assumes that the paired comparisons are generated according to a generalized linear model, where paired comparisons are independent, and item i is preferred to item j with probability

$$P_{ij} = H(\nu_i - \nu_j)$$

where $\nu \in \mathbb{R}^n$ is a vector denoting the strength, rank, or skill level of the n players. In the context of the GLM model, Fogel et al. propose the following similarity matrix

$$S_{i,j}^{glm} = \sum_{k=1}^n \mathbf{1}_{\{m_{i,k}m_{j,k} > 0\}} \left(1 - \frac{|C_{i,k} - C_{j,k}|}{2} \right) + \frac{\mathbf{1}_{\{m_{i,k}m_{j,k} = 0\}}}{2} \quad (44)$$

where $m_{i,k} = 1$ if i and j played in a match, and 0 otherwise. We denote by SER-GLM the Serial-Rank algorithm based on the above GLM model given by (44). For the cardinal case, SER-GLM is roughly similar to SER, except in the complete graph case, when SER-GLM consistently outperforms SER under both the MUN and ERO noise models. We refer the reader to Figures 5 and 6 for the numerical results showing how SER and SER-GLM compare to the existing and newly proposed methods.

APPENDIX C RANK CENTRALITY

In recent work [27], Negahban et al. propose an iterative algorithm for the rank aggregation problem by estimating scores for the items from the stationary distribution of a certain random walk on the graph of items, where edges encode the outcome of pairwise comparisons. The authors propose this approach in the context of the rank aggregation problem, which, given as input a collection of sets of pairwise comparisons over n players (where each set is provided by an independent rating system) the goal is to provide a global ranking that is as consistent as possible with the given measurements of all k ranking systems. At each iteration of the random walk, the probability of transitioning from vertex i to vertex j is directly proportional to how often player j beat player i across all the matches the two players confronted, and is zero if the two players have never played a match before. In other words, the random walk has a higher chance of transitioning to a more skillful neighbors, and thus the frequency of visiting a particular node, which reflects the rank or the skill level of the corresponding players, is thus encoded in the stationary distribution of the associated Markov Chain. The authors of [27] propose the following approach for computing the Markov matrix, which we adjust to render it applicable to both ordinal and cardinal measurements, in the case of a single rating system. Note that in Section 6 where we discuss the rank aggregation problem in the context of multiple rating systems, we rely on the original version of the algorithm. For a pair of items i and j , let $Y_{ij}^{(l)}$ be equal to 1 if player j beats player i , and 0 otherwise, during the l^{th} match between the two players, with $l = 1, \dots, k$. The BTL model assumes that $\mathbb{P}(Y_{ij}^{(l)}) = \frac{w_j}{w_i + w_j}$, where w represent the underlying vector of positive real weights associated to each player. The approach in [27] starts by estimating the fraction of times players j has defeated player i , which is denoted by

$$a_{ij} = \frac{1}{k} \sum_{l=1}^k Y_{ij}^{(l)}, \quad (45)$$

as long as players i and j competed in at least one match, and zero otherwise. Next, consider the symmetric matrix

$$A_{ij} = \frac{a_{ij}}{a_{ij} + a_{ji}}, \quad (46)$$

which converges to $\frac{w_j}{w_i + w_j}$, as $k \rightarrow \infty$. To define a valid transition probability matrix, the authors of [27] scale all the edge weights by $1/d_{max}$ and consider the resulting random walk

$$P_{ij} = \begin{cases} \frac{1}{d_{max}} A_{ij} & \text{if } i \neq j \\ 1 - \frac{1}{d_{max}} \sum_{k \neq i} A_{ik} & \text{if } i = j, \end{cases} \quad (47)$$

where d_{max} denotes the maximum out-degree of a node, thus making sure that each row sums to 1. The stationary distribution π is the top left eigenvector of P , and its entries denote the final numerical scores associated to each node, which, upon sorting, induce a ranking of the n players.

To render the above approach applicable⁹ in the case

9. Otherwise, in the ordinal case, A_{ij} is either 0 or 1

when $k = 1$ of a single rating system (for both cardinal and ordinal measurement), but also when $k > 1$ for the case of cardinal measurements, we propose the following alternatives to designing the winning probabilities a_{ij} given by (45), and inherently the final winning probability matrix A in (46). Once we have an estimate for A (given by (50) for ordinal data, respectively by (51) for cardinal data), we proceed with building the transition probability P as in (47) and consider its stationary distribution. Note that we also make these two new methods proposed below applicable to the setting of multiple rating systems $k > 1$, by simply averaging out the resulting winning probabilities $A_{ij}^{(l)}$, $l = 1, \dots, k$, given by each rating system via (50) or (51), across all rating systems

$$A_{ij} = \frac{1}{k} \sum_{l=1}^k A_{ij}^{(l)}, \quad (48)$$

and then consider the transition probability matrix P as in (47) and its stationary distribution.

C.1 Adjusted Rank-Centrality for ordinal measurements

To handle the case of ordinal measurements, we propose a hybrid approach that combines Serial-Rank and Rank-Centrality, and yields more accurate results than the former one, as it can be seen in Figure 6. We proceed by computing the S^{match} matrix as in the Serial-Rank algorithm (given by (41) and (42) in Appendix B) that counts the number of matching comparisons between i and j with other third reference items k . The intuition behind this similarity measure is that players that beat the same players and are beaten by the same players should have a similar ranking in the final solution. Note that $S_{ij} \leq n$ for any pair of players¹⁰, and thus $\frac{S_{ij}}{n} \in [0, 1]$

Note that, whenever S_{ij} is very large, say $S_{ij} \approx n$, meaning the two players are very similar, then the quantity $1 - \frac{S_{ij}}{n}$ is small and close to zero, and thus a good proxy for the difference in the winning probabilities A_{ij} and A_{ji} defined in (46). In other words, if two players are very similar, it is unlikely that, had they played a lot of matches against each other, one player will defeat the other in most of the matches. On the other hand, if two players are very dissimilar, and thus S_{ij} is close to zero and $1 - \frac{S_{ij}}{n}$ is close to one, then it must be that, had the two players met in many matches, one would have defeated the other in a significant fraction of the games. With these observations in mind, and in the spirit of (46), we design the matrix A of winning probabilities such that

$$\begin{cases} A_{ij} + A_{ji} = 1 \\ |A_{ij} - A_{ji}| = 1 - \frac{S_{ij}}{n}. \end{cases} \quad (49)$$

for a pair of players that met in a game. We lean the balance in favor of the player who won in the (single) direct match, and assign to him the larger winning probability. Keeping in mind that A_{ij} should be a proxy for the fraction of times

player j defeated player i (thus whenever $C_{ij} > 0$ it must be that $A_{ij} > A_{ji}$), the above system of equations (49) yields

$$A_{ij} = \begin{cases} 1 - \frac{1}{2} \frac{S_{ij}}{n}, & \text{if } C_{ij} > 0 \\ \frac{1}{2} \frac{S_{ij}}{n}, & \text{if } C_{ij} < 0. \\ 0, & \text{if } C_{ij} = 0. \end{cases} \quad (50)$$

We remark that, in the case of outliers given by the ERO noise model (29), our above proposed version of Rank-Centrality (denoted as RC), when used in the setting of multiple rating systems, performs much better than the original Rank-Centrality algorithm (denoted as RCO), as shown in the bottom plot of Figure 9.

C.2 Adjusted Rank-Centrality for cardinal measurements

For the case of cardinal measurements, we propose a similar matrix A of winning probabilities, and incorporate the magnitude of the score into the entries of A . The intuition behind defining the winning probability is given by the following reasoning. Whenever C_{ij} takes the largest possible (absolute) value (i.e., assume $C_{ij} = (n-1)$, thus j defeats i by far), we define the winning probability that player j defeats player i to be largest possible, i.e., $A_{ij} = 1$, and in general, the larger the magnitude of C_{ij} , the larger A_{ij} should be. On the other hand, whenever C_{ij} has the smallest possible (absolute) value (i.e., assume $C_{ij} = 1$), then the winning probability should be as small as possible, i.e., close to $\frac{1}{2}$. With these two observations in mind, we define the winning probability matrix as

$$A_{ij} = \begin{cases} \frac{1}{2} + \frac{1}{2} \frac{C_{ij}}{n-1}, & \text{if } C_{ij} > 0 \\ \frac{1}{2} - \frac{1}{2} \frac{C_{ij}}{n-1}, & \text{if } C_{ij} < 0. \\ 0, & \text{if } C_{ij} = 0. \end{cases} \quad (51)$$

APPENDIX D

FUTURE RESEARCH DIRECTIONS

We highlight in this section several possible future research directions that we believe are interesting and worth pursuing further.

D.1 Robustness to noise of SVD-based ranking

An interesting research direction is the analysis of the robustness to noise of the SVD-based ranking method discussed in Section 2.3. We remark here that, for the **Erdős-Rényi Outliers** $\text{ERO}(n, p, \eta)$ model (29), the following decomposition could render the SVD-Rank method amenable to a theoretical analysis using tools from the random matrix theory literature on rank-2 deformations of random matrices [6]. Note that the expected value of the entries of C is given by

$$\mathbb{E}C_{ij} = (r_i - r_j)(1 - \eta)p, \quad (52)$$

in other words, $\mathbb{E}C$ is a rank-2 skew-symmetric matrix

$$\mathbb{E}C = (1 - \eta)p(re^T - er^T) \quad (53)$$

Next, one may decompose the given data matrix C as

$$C = \mathbb{E}C + R \quad (54)$$

10. This is since C_{ii} is defined to be 1, otherwise it would be true that $S_{ij} \leq n - 2$.

where $R = C - \mathbb{E}C$ is a random skew-symmetric matrix whose elements have zero mean and are given by

$$R_{ij} = \begin{cases} (r_i - r_j)[1 - (1 - \eta)p] & \text{w.p. } (1 - \eta)p \\ q - (r_i - r_j)(1 - \eta)p & \text{w.p. } \eta p \\ -(r_i - r_j)(1 - \eta)p & \text{w.p. } 1 - p, \end{cases} \quad (55)$$

whenever $i \neq j$ and where $q \sim \text{Unif.}[-(n-1), n-1]$, which renders the given data matrix C decomposable into a low-rank (rank-2) perturbation of a random skew-symmetric matrix. The case $p = 1$ which corresponds to the complete graph $G = K_n$ simplifies (55), and is perhaps a first step towards a theoretical investigation.

D.2 Ranking with soft (edge) constraints as a generalized eigenvalue problem

We have seen in the previous Section 7 that adding hard constraints to the ranking problem improves the end results. A natural question to consider is how one may incorporate soft constraints into the problem formulation. Let H denote the Hermitian matrix of available pairwise cardinal rank-offsets after embedding in the angular space, and similarly let Q denote the Hermitian matrix of pairwise cardinal rank-offset *constraints* the user wishes to enforce, also after mapping to the angular embedding space. One possibility is to incorporate weights into either the spectral or SDP relaxation of the synchronization problem, in which case, the input to the synchronization formulation will be given by

$$\tilde{H} = H + \lambda Q$$

where Q is the matrix of rank-offset (soft) constraints. The larger the parameter λ , the more weight is given to the available pairwise constraints, and thus the more likely they are to be satisfied in the final proposed solution. However, in certain instances one may wish to enforce a lower bound on how well the given constraints are satisfied, in which case we adjust the optimization problem as follows. We maximize the same quadratic form $z^* H z$ as in (12), but subject to the condition that the given constraints are satisfied well enough, i.e., $z^* Q z \geq \alpha$, for a user specified α , and that the sum of the absolute values of all complex numbers z_i is n

$$\begin{aligned} & \max_{z=(z_1, \dots, z_l) \in \mathbb{C}^n} && z^* H z \\ & \text{subject to} && z^* Q z \geq \alpha \\ & && z^* z = n. \end{aligned} \quad (56)$$

We remark that this approach is extremely similar to the one considered in the constrained clustering algorithm proposed in [34], where the goal is to cluster a given weighted graph with adjacency matrix H , subject to soft constraints captured in a sparse matrix Q , of the form $Q_{ij} = 1$ (respectively, $Q_{ij} = -1$) if nodes i and j should be (respectively, should not be) in the same cluster, and $Q_{ij} = 0$ if no information is available. As shown in [34], one arrives at a generalized eigenvalue problem based on the matrices H and Q . However, instead of working with the above H and Q matrices (which in our setting (56) are Hermitian matrices) one could instead consider their corresponding graph Connection Laplacians, and solve the resulting generalized eigenvalue problem where both matrices are now symmetric diagonally dominant. We point out our recent

work [11], where we present a principled spectral approach to the above well-studied constrained clustering problem, and reduce it to a generalized eigenvalue problem involving Laplacian matrices that can be solved in nearly-linear time due to recent progress in the area of Laplacian linear systems solvers. In practice, this translates to a very fast implementation that consistently outperforms the approach in [34], both in terms of accuracy and running time. It would be interesting to explore whether our approach in [11] can be extended to the setting of constrained ranking.

D.3 (Signless) Ranking with Unsigned Cardinal Comparisons

Consider for example the scenario where one is given the score sheet of all soccer games in the England Football Premier League, recording the goal difference for each game, but without disclosing who won and who lost the game. In other words, one wishes to reconstruct the underlying rankings (up to a global reversal of the ordering) only using information on the magnitude of the rank offsets

$$C_{ij} = |r_i - r_j|. \quad (57)$$

We remark that the resulting problem is nothing but an instance of the *graph realization problem* on the line, with noisy distances [12]. In practice, an instance of this problem arises in shotgun genome sequencing experiments, where the goal is to reorder substrings of cloned DNA strands (called *reads*) using assembly algorithms that exploit the overlap between the reads. We refer the reader to the recent work of Fogel et al. on convex relaxations for the *seriation problem* [19], including an application to the above shotgun gene sequencing task. We remark that preliminary numerical results indicate that SVD-based ranking performs remarkably well in this setting of recovering the ordering (up to a global sign) using information only on the magnitude of the rank offsets but not their sign, and defer this investigation to future work.

D.4 Ranking both items and raters

Another possible direction concerns the detection of outlier voters, in the setting when there are multiple voting systems providing pairwise comparisons on the same set of players. For example, in Section 6 we considered the problem of rank aggregation, of combining information from multiple voters or rating systems, which provide incomplete and inconsistent rankings or pairwise comparisons between the same set of players or items. A natural question asks whether it is possible to derive a reputation for both raters and items, and identify outlier voters of low credibility. Furthermore, one can incorporate the credibility of the voters when estimating the reputation of the objects, as considered in the recent work of de Kerchove and Van Dooren [14], who introduce a class of iterative voting systems that assign a reputation to both voters and items, by applying a filter to the votes. Another potential application is to crowdsourcing [23], where crowd workers have different levels of competencies and are likely to produce noisy judgments regarding ranking of items.

D.5 Ranking via l_1 Angular Synchronization

Singer [30] and Yu [38] observe that embedding in the angular space is significantly more robust to outliers when compared to l_1 or l_2 based embeddings in the linear space. Since the existing spectral and SDP relaxations for the angular synchronization problem [30] approximate a least squares solution in the angular space, a natural question to ask is whether an l_1 -based formulation of the angular synchronization problem could further improve the robustness to noise or sparsity of the measurement graph. A preliminary investigation of such an l_1 based approach for angular synchronization (which we did not further apply to the ranking problem) using the splitting orthogonality constraints approach of Lai and Osher [24] based on Bregman iterations, suggests that it often yields more accurate results than the existing spectral relaxation of its least-squares solution only for favorable noise regimes. At high levels of noise, the spectral relaxation performs better, and does so at a much lower computational cost. On a related note, we point out the SDP relaxation of the Least Unsquared Deviation (LUD) formulation [33], for which the authors give exact recovery guarantees under the Erdős-Rényi Outliers model.

D.6 Ranking over the Sphere and Synchronization over SO(3)

Another research direction we find interesting concerns the possibility of casting the ranking problem as a synchronization problem over the group SO(3), as opposed to SO(2) which we considered in this paper. Consider for example a sphere centered at the origin, with points on the sphere corresponding to the players and the z -axis to their ranks, such that the closer a player is to the South Pole, the higher her or his ranking is. The extra *dimension* would facilitate the extraction of partial rankings, which one can collect by traversing along longitudinal lines and their immediate neighborhoods. In addition, one can also extract sets of players which are about the same rank (roughly speaking, they have about the same latitude), but cannot be positioned accurately with respect to each other (for example, a set of 10 players out of 100, are ranked in positions 81 – 90, but based on the data, their relative ordering cannot be established). In other words, if one considers a slice of the sphere (parallel to the xy -plane) all the players on or nearby it would have about the same rank. Perhaps most important, it would also be interesting to investigate the extent to which the extra dimension increases the overall robustness to noise of the algorithm. Note that, while for SO(2), finding an optimal circular permutation was required to eliminate the additional degree of freedom, for the case of SO(3) one could search for the best rotation around the origin of the sphere, such that the resulting rank offsets agree with the initial data as best as possible.

D.7 Robustness to noise

Another future direction would be to aim for guarantees in terms of robustness to noise. We remark that existing theoretical guarantees from the recent literature on the group synchronization problem by various works of Bandeira, Boumal, Charikar, Singer and Spielman [4], [5], [30], trivially

translate to lower bounds on the largest amount of noise permissible in the cardinal comparisons for an Erdős-Rényi outliers noise model, while still achieving an approximate solution that correlates well with the underlying ground truth ranking. However, note that a perfect or close-to-perfect recovery of the angles in the angular synchronization problem is not a necessary condition for a similar recovery of the underlying ground truth ranking, since it is enough that only the relative ordering of the angles is preserved. Furthermore, the main theorem of [4] provides an l_∞ error on the MLE estimator for angular synchronization, which, when compared to the minimal distance between two consecutive angles, could lead to theoretical bounds for exact recovery of the rankings¹¹.

Another interesting possible direction is to investigate the robustness to noise of Sync-Rank in the setting of the, perhaps more realistic, multiplicative uniform noise model (28) and explain the empirical findings in Figure 4 (a) and (b). Finally, the surprising performance (relative to other methods) of Sync-Rank based on rank offsets given by (24) and illustrated in Figure 8 of Section 5.4 across 30 years of data, raises the question of how does a disc measurement graph affect the performance of the algorithms, i.e., in the (often very realistic) setting where nodes (players) lie on a one-dimensional line, and there exists an edge (i.e., match) between two nodes if and only if they are at most r units (ranks) apart.

D.8 The Minimum Linear Arrangement Problem

One could also investigate any potential connections between synchronization and SVD-based ranking using only magnitude information (as detailed in Section D.3) and the Minimum Linear Arrangement (MLA) problem, which is defined as follows. Given a graph $G = (V, E)$ and positive edge weights $w : E \mapsto \mathbb{R}_+$, a linear arrangement is a permutation $\pi : V \mapsto \{1, 2, \dots, n\}$. The cost of the arrangement is given by

$$\sum_{ij \in E} w(u, v) |\pi(u) - \pi(v)| \quad (58)$$

In the MLA problem, the goal is to find a linear arrangement of minimum cost a task known to be NP-complete. In the approximation theory literature, the work of Feige and Lee [15] provides an $O(\sqrt{\log n \log \log n})$ -approximation SDP-based algorithm for the MLA problem.

APPENDIX E ENGLISH PREMIER LEAGUE STANDINGS

In Table 2 we show the rankings obtained by each algorithm for the English Premier League 2013-2014 data set, based on the input matrix C^{nw} which counts the number of net wins between a pair of teams. We also present the final official ranking at the end of the season, which are denoted by GT.

11. We thank one of the referees for pointing out this connection.

Team	SVD	LS	SER	SYNC	SER-GLM	SYNC SUP	SYNC SDP	RC	GT
Arsenal	1	4	5	4	4	5	4	4	4
Aston Villa	19	16	18	16	20	19	16	16	15
Cardiff	17	20	16	20	15	20	20	20	20
Chelsea	8	1	1	1	2	1	1	1	3
Crystal Palace	14	12	14	12	16	13	12	13	11
Everton	5	5	4	5	5	6	5	5	5
Fulham	15	17	17	17	17	16	17	17	19
Hull	13	14	13	14	13	14	14	14	16
Liverpool	3	3	3	3	3	3	3	3	2
Man City	2	2	2	2	1	2	2	2	1
Man United	4	7	7	7	8	7	7	7	7
Newcastle	10	11	10	11	11	10	11	11	10
Norwich	18	15	19	15	18	18	15	15	18
Southampton	7	8	8	8	7	8	8	8	8
Stoke	11	9	9	9	9	9	9	9	9
Sunderland	20	18	20	18	19	17	18	18	14
Swansea	9	10	11	10	10	12	10	10	12
Tottenham	6	6	6	6	6	4	6	6	6
West Brom	16	19	15	19	12	15	19	19	17
West Ham	12	13	12	13	14	11	13	12	13
Nr. upsets $Q^{(u)}$	66	44	44	44	52	48	44	46	54
Corr w. GT	0.69	0.87	0.80	0.87	0.75	0.84	0.87	0.86	1.00

TABLE 2

English Premier League Standings 2013-2014, based on the input matrix C^{nw} (which counts the number of net wins between a pair of teams). GT denotes the final official ranking at the end of the season.

The bottom row shows the correlation of the result obtained by each algorithm with the ground truth ranking ¹².

12. We do not believe that correlation with the official standings is a good measure of success, simply because they are based on different rules, i.e., on accumulation of points for each win, tie, or loss. One can think of the four different types of pre-processing criteria as different input data sets of pairwise comparisons, and the goal is to propose a solution that best agrees with the input data, whatever that is.