Graph Resistance and Learning from Pairwise Comparisons

Julien M. Hendrickx[*,†] **Alex Olshevsky**[†] **Venkatesh Saligrama**[†]

Abstract

We consider the problem of learning the qualities of a collection of items by performing noisy comparisons among them. Following the standard paradigm, we assume there is a fixed "comparison graph" and every neighboring pair of items in this graph is compared k times according to the Bradley-Terry-Luce model (where the probability than an item wins a comparison is proportional the item quality). We are interested in how the relative error in quality estimation scales with the comparison graph in the regime where k is large. We prove that, after a known transition period, the relevant graph-theoretic quantity is the square root of the resistance of the comparison graph. Specifically, we provide an algorithm that is minimax optimal. The algorithm has a relative error decay that scales with the square root of the graph resistance, and provide a matching lower bound (up to log factors). The performance guarantee of our algorithm, both in terms of the graph and the skewness of the item quality distribution, outperforms earlier results.

1. Introduction

This paper considers quality estimation from pairwise comparisons, which is a common method of preference elicitation from users. For example, the preference of a customer for one product over another can be thought of as the outcome of a comparison. Because customers are idiosyncratic, such outcomes will be noisy functions of the quality of the underlying items. A similar problem arises in crowdsourcing systems, which must strive for accurate inference even in the presence of unreliable or error-prone participants. Because crowdsourced tasks pay relatively little, errors are common; even among workers making a genuine effort, inherent ambiguity in the task might lead to some randomness in the outcome. These considerations make the underlying estimation algorithm an important part of any crowdsourcing scheme.

Our goal is accurate inference of true item quality from a collection of outcomes of noisy comparisons. We will use one of the simplest parametric models for the outcome of comparisons, the Bradley-Terry-Luce (BTL) model, which associates a real-valued quality measure to each item and posits that customers select an item with a probability that is proportional to its quality. *Given a "comparison graph" which captures which pairs of items are to be compared, our goal is to understand how accuracy scales in terms of this graph when participants make choices according to the BTL model.*

We focus on the regime where we perform many comparisons of each pair of items in the graph. In this regime, we are able to give a satisfactory answer to the underlying question. Informally, we prove that, up to various constants and logarithms, the relative estimation error will scale with the square root of measures of resistance in the underlying graph. Specifically, we propose an algorithm whose performance scales with graph resistance, as well as a matching lower bound. The difference between our upper and lower bounds depends only on the log of the confidence level and on the skewness of the item qualities. Additionally, we note that our performance guarantees scale better in terms of item skewness as compared to previous work.

1.1. Formal problem statement

We are given an undirected "comparison graph" $G(V,E)$, where each node i has a positive weight w_i. If $(i,j) \in E$, then we perform k comparisons between i and j. The outcomes of these comparisons are i.i.d. Bernoulli and the probability that i wins a given comparison according to the

[*]Equal contribution . Correspondence to: Alex Olshevsky <alexols@bu.edu>.

Proceedings of the 36^{th} International Conference on Machine Learning, Long Beach, California, PMLR 97, 2019. Copyright 2019 by the author(s).

[*] Department of Mathematical Engineering, ICTEAM, UCLouvain, Belgium
[†] Department of Electrical and Computer Engineering, Boston University, USA
This work was supported by the National Science Foundation Grant CCF 1527618, Office of Naval Research Grant N0014-18-1-2257, the U.S. Department of Homeland Security, Science andTechnology Directorate, Office of University Programs under Grant 2013-ST-061-ED0001, Communauté française de Belgique - Actions de Recherche Concertées," and a WBI World excellence fellowship.

BTL model is

$$p_{ij} = \frac{w_i}{w_i + w_j} \quad (1)$$

The goal is to recover the weights w_i from the outcomes of these comparisons. Because multiplying all w_i by the same constant does not affect the distribution of outcomes, we will recover a scaled version of the weight vector w.

Thus our goal will thus be come up with a vector of estimated weights $\hat{W}$ close, in a scale-invariant sense, to the true but unknown vector[1] w. A natural error measure turns out to be the absolute value of the sine of the angle defined by w and $\hat{W}$, which can also be expressed as (see Lemma **??** in the Supplementary Information)

$$\left|\sin(\hat{W}, w)\right| = \inf_{\alpha \in \mathbb{R}} \frac{||\hat{W} - \alpha w||_2}{||\alpha w||_2}. \quad (2)$$

In other words, $|\sin(\hat{W}, w)|$ is the relative error to the closest normalization of the true quality vector w. We will also discuss the connection between this error measure and others later on in the paper.

Following earlier literature, we assume that

$$\max_{i,j \in V} \frac{w_i}{w_j} \quad \leq \quad b$$

for some constant b. The number b can be thought of as a measure of the skewness of the underlying item quality. Our goal is to understand how the error between $\hat{W}$ and w scales as a function of the comparison graph G.

1.2. Literature Review

The dominant approach to recommendation systems relies on inferring item quality from raw scores provided by users (see (Jannach et al., 2016)). However, such scores might be poorly calibrated and inconsistent; alternative approaches that offer simpler choices might perform better.

Our starting point is the Bradley-Terry-Luce (BTL) model of Eq. (1), dating back to (Bradley & Terry, 1952; Luce, 2012), which models how individuals make noisy choices between items. A number of other models in the literature have also been used as the basis of inference, we mention the Mallows model introduced in (Mallows, 1957) and the PL and Thurstone models (see description in (Hajek et al., 2014)). However, we focus here solely on the BTL model.

Our work is most closely related to the papers (Negahban et al., 2012) and (Negahban et al., 2016). These works proposed an eigenvector calculation which, provided the number of comparisons is sufficiently large, successfully recovers the true weights w from the outcomes of noisy comparisons. The main result of (Negahban et al., 2016) stated that, given a comparison graph, if the number of comparisons per edge satisfied a certain lower bound, then it is possible to construct an estimate $\hat{W}$ satisfying

$$\frac{||\hat{W} - w||_2}{||w||_2} \leq O\left(\frac{b^{5/2} d_{\max}}{d_{\min}(1-\lambda)} \sqrt{\frac{\log n}{k d_{\max}}}\right) \quad (3)$$

with high probability, where $d_{\min}, d_{\max}$ are, respectively, the smallest and largest degrees in the comparison graph, $1 - \lambda$ is the spectral gap of a certain normalized Laplacian of the comparison graph, and both $w, \hat{W}$ are normalized so that their entries sum to 1. It can be proved (see Lemma **??**) that the relative error on the left-hand side of Eq. (3) is within a $\sqrt{b}$ factor of the measure $|\sin(\hat{W}, w)|$ provided that $\max_{i,j} \hat{W}_i/\hat{W}_j \leq b$, so asymptotically these two measures differ only by factor depending on the skewness b.

The problem of recovering w was further studied in (Rajkumar & Agarwal, 2014), where the comparison graph was taken to be a complete graph but with comparisons on edges made at non-uniform rates. The sample complexity of recovering the true weights was provided as a function of the smallest sampling rate over pairs of items.

A somewhat more general setting was considered in (Shah et al., 2016), which considered a wider class of noisy comparison models which include the BTL model as a special case. Upper and lower bounds on the minimax optimal rates in estimation, depending on the eigenvalues of a corresponding Laplacian, were obtained for absolute error in several different metrics; in one of these metric, the Laplacian semi-metric, the upper and lower bounds were tight up to constant factors. Similarly to (Shah et al., 2016), our goal is to understand the dependence on the underlying graph, albeit in the simpler setting of the BTL model.

Our approach to the problem very closely parallels the approach of (Jiang et al., 2011), where a collection of potentially inconsistent rankings is optimally reconciled by solving an optimization problem over the comparison graph. However, whereas (Jiang et al., 2011) solves a linear programming problem, we will use a linear least squares approach, after a certain logarithmic change of variable.

We now move on to discuss work more distantly related to the present paper. We mention that the problem we study here is related, but not identical, to the so-called noisy sorting problem, introduced in (Braverman & Mossel, 2009), where better items win with probability at least $1/2 + \delta$ for some positive δ. This assumption does not hold for the BTL model with arbitrary weights. Noisy sorting was also studied in the more general setting of ranking models satisfying a transitivity condition in (Shah et al., 2017) and (Pananjady et al., 2017), where near-optimal minimax rates were derived. Finally, optimal minimax rates for noisy sorting were recently demonstrated in (Mao et al., 2017).

[1] We follow the usual convention of denoting random variables by capital letters, which is why $\hat{W}$ is capitalized while w is not.

There are a number of variations of this problem that have been studied in the literature which we do not survey at length due to space constraints. For example, the papers (Yue et al., 2012; Szörényi et al., 2015) considered the online version of this problem with corresponding regret, (Chen & Suh, 2015) considered recovering the top K ranked items, (Falahatgar et al., 2017; Agarwal et al., 2017; Maystre & Grossglauser, 2015) consider recovering a ranked list of the items, and (Ajtai et al., 2016) consider a model where comparisons are not noisy if the item qualities are sufficiently far apart. We refer the reader to the references within those papers for more details on related works in these directions.

1.3. Our approach

We will construct our estimate $\hat{W}$ by solving a log-least-squares problem described next. We denote by F_{ij} the fraction of times node i wins the comparison against its neighbor j, and we further set $R_{ij} = F_{ij}/F_{ji}$. As the number of comparisons on each edge goes to infinity, we will have that R_{ij} approaches w_i/w_j with probability one. Our method consists in finding $\hat{W}$ as follows:

$$\hat{W} = \arg\min_{v\in\mathbb{R}_+^{|E|}} \sum_{(i,j)\in E} (\log(v_i/v_j) - \log R_{ij})^2 \qquad (4)$$

This can be done efficiently by observing that it amounts to solving the linear system of equations

$$\log \hat{W}_i - \log \hat{W}_j = \log R_{ij}, \quad \text{for all } (i,j) \in E,$$

in the least square sense. Let B to be the incidence matrix[2] of the comparison graph. Stacking up the R_{ij} into a vector R, we can then write

$$B^T \log \hat{W} = \log R$$

Least-square solutions satisfy

$$BB^T \log \hat{W} = B \log R$$

or equivalently $L \log \hat{W} = B \log R$, where $L = BB^T$ is the graph Laplacian. Finally, a solution is given by

$$\log \hat{W} = L^{\dagger} B \log R. \qquad (5)$$

where $L^{\dagger}$ is the Moore-Penrose pseudoinverse. By using the classic results of (Spielman & Teng, 2014), Eq. (5) can be solved for $\hat{W}$ to accuracy ϵ in nearly linear time in terms of the size of the input, specifically in $O(|E| \log^c n \log(1/\epsilon))$ iterations for some constant $c > 0$. We note that, for connected graphs, all solutions w of (4) are equal up to a multiplicative constant and are thus equivalent in terms of criterion (2).

[2] Given an directed graph with n nodes and $|E|$ edges, the incidence matrix is the $n \times |E|$ matrix whose i'th column has a 1 corresponding to the source of edge i, a -1 corresponding to the destination of node i, and zeros elsewhere. For an undirected graph, an incidence matrix is obtained by first orienting the edges arbitrarily.

1.4. Our contribution

We will find it useful to view the graph as a circuit with a unit resistor on each edge; Ω_{ij} will denote the resistance between nodes i and j in this circuit, $\Omega_{\max}$ denotes the largest of these resistances over all pairs of nodes $i, j = 1, \ldots, n$ and similarly Ω_{avg} denotes the average resistance over all pairs. We will use E_{ij} to denote the set of edges lying on at least one simple path starting at i and terminating at j, with $E_{\max}$ denoting the largest of the E_{ij}. Naturally, $E_{\max}$ is upper bounded by the total number of edges in the comparison graph. The performance of our algorithms is described by the following theorem.

Theorem 1. *Let $\delta \in (0, e^{-1})$. There exist absolute constants constants c_1, c_2 such that, if $C_{n,\delta} \geq c_1 \log(n/\delta)$ and $k \geq c_2 E_{\max} C_{n,\delta}^2$ and $k \geq c_3 \Omega b^2 (1 + (\log(1/\delta))$, then we have, with probability at least $1 - \delta$, that*

$$\sin(\hat{W}, w)^2 \leq O\left(\frac{\min\left(b^2\Omega_{\max}, b^4\Omega_{\text{avg}}\right)}{k} \times \left(\left(1 + \log\frac{1}{\delta}\right) + \frac{E_{\max} C_{n,\delta}^2}{k} \right) \right)$$

The main feature of this theorem is the favorable form of the bound in the setting when k is large. Then only the leading term

$$\frac{\min(b^2\Omega_{\max}, b^4\Omega_{\text{avg}})(1 + \log 1/\delta)}{k}$$

dominates the expression on the right-hand-side. Taking square roots, it follows that, asymptotically,

$$\left|\sin(\hat{W}, w)\right| = \widetilde{O}\left(\sqrt{\frac{b^2\Omega_{\max}}{k}}\right) \text{ and } \widetilde{O}\left(\sqrt{\frac{b^4\Omega_{avg}}{k}}\right),$$

where the $\widetilde{O}$ notation hides logarithmic factor in δ.

Our other main result is that, in the regime when k is large, there is very little room for improvement.

Theorem 2. *For any comparison graph G, and for any algorithm, as long as $k \geq c\sqrt{\lambda_{\max}(L)} n \Omega_{\text{avg}}$ for some absolute constant c, we have that*

$$\sup_{w\in\mathbb{R}_+^n} E\left|\sin(\hat{W}, w)\right| \geq \Omega\left(\sqrt{\frac{\Omega_{avg}}{k}}\right),$$

where as before L is the graph Laplacian.

Comparing Theorem 1 with Theorem 2, we see that the performance bounds of Theorem 1 are minimax optimal, at least up to the logarithmic factor in the confidence level δ and dependence on the skewness factor b. *We can thus conclude that the square root of the graph resistance is the key graph-theoretic property which captures how relative error decays for learning from pairwise comparisons.* This observation is the main contribution of this paper.

Table 1. Comparison, for different families of graphs, of $\widetilde{O}\left(\frac{d_{\max}}{d_{\min}(1-\lambda)}\sqrt{\frac{1}{d_{\max}}}\right)$ and $\widetilde{O}(\sqrt{b\Omega_{\max}})$, which are, respectively, the asymptotic bounds (3) in (Negahban et al., 2016), and the first bound from our Theorem 1. The common decay in $k^{-1/2}$ is omitted for the sake of conciseness.

Graph	Eq. (3)	Theorem 1
Line	$b^{5/2}n^2$	$b\sqrt{n}$
Circle	$b^{5/2}n^2$	$b\sqrt{n}$
2D grid	$b^{5/2}n$	b
3D grid	$b^{5/2}n^{2/3}$	b
Star graph	$b^{5/2}\sqrt{n}$	b
2 stars joined at centers	$b^{5/2}n^{1.5}$	b
Barbell graph	$b^{5/2}n^{3.5}$	$b\sqrt{n}$
Geo. random graph	$b^{5/2}n$	b
Erdos-Renyi	$b^{5/2}$	b

1.5. Comparison to previous work

Table 1 quantifies how much the bound of Theorem 1 expressed in terms of $\Omega_{\max}$ improves the asymptotic decay rate on various graphs over the bound (Negahban et al., 2016). The $\widetilde{O}$ notation ignores log-factors. Both random graphs are taken at a constant multiple threshold which guarantees connectivity; for Erdos-Renyi this means $p = O((\log n)/n)$ and for a geometric random graph, this means connecting nodes at random positions at the unit square when they are $O\left(\sqrt{(\log n)/n}\right)$ apart.

Most of the scalings for eigenvalues of normalized Laplacians used in Table 1 are either known or easy to derive. For an analysis of the eigenvalue of the barbell graph[3], we refer the reader to (Landau & Odlyzko, 1981); for mixing times on the geometric random graph, we refer the reader to (Avin & Ercal, 2007); for the resistance of an Erdos-Renyi graph, we refer the reader to (Sylvester, 2016).

In terms of the worst-case performance in terms of the number of nodes, our bound grows at worst as $\widetilde{O}\left(b\sqrt{n/k}\right)$ using the observation that $\Omega_{\max} = O(n)$. By contrast, for the barbell graph, the bound of (Negahban et al., 2016) grows as $\widetilde{O}(b^{5/2}n^{3.5}/\sqrt{k})$, and it is not hard to see this is actually the worst-case scaling in terms of the number of nodes.

Finally, we note that these comparisons use slightly different error measures: $|\sin(\hat{W}, w)|$ on our end vs the relative error in the 2-norm after $w, \hat{W}$ have been normalized to sum to one, used by (Negahban et al., 2016). To compare both in terms of the latter, we could multiply our bounds by $\sqrt{b}$ (see Lemma **??**).

[3]Following (Wilf, 1989), the barbell graph refers to two complete graphs on $n/3$ vertices connected by a line of $n/3$ vertices.

1.6. Notation

The remainder of this paper is dedicated to the proof Theorem 1 (Theorem 2 is proved in the Supplementary Information). However, we first collect some notation we will find occasion to use.

As mentioned earlier, we let F_{ij} be the empirical rate of success of item i in the k comparisons between i and j; thus $E[F_{ij}] = p_{ij}$ so that the previously introduced R_{ij} can be expressed as $R_{ij} = \frac{F_{ij}}{F_{ji}}$. We also let $\rho_{ij} = w_i/w_j = p_{ij}/p_{ji}$, to which R_{ij} should converge asymptotically.

We will make a habit of stacking any of the quantities defined into vectors; thus F, for example, denotes the vector in $\mathbb{R}^{|E|}$ which stacks up the quantities F_{ij} with the choice of i and j consistent with the orientation in the incidence matrix B. The the vectors p and ρ are defined likewise.

2. Proof of the algorithm performance (Theorem 1)

We begin the proof with a sequence of lemmas which work their way to the main theorem. The first step is to introduce some notation for the comparison on the edge (i, j).

Let X_{ij} be the outcome of a single coin toss comparing coins i and j. Using the standard formula for the variance of a Bernoulli random variable, we obtain

$$\begin{aligned} \mathrm{Var}(X_{ij}) &= p_{ij}(1-p_{ij}) = \frac{w_i w_j}{(w_i+w_j)^2} \\ &= \frac{1}{\rho_{ij}+2+\rho_{ij}^{-1}} =: \frac{1}{v_{ij}}, \end{aligned} \tag{6}$$

where we have defined $v_{ij} = \rho_{ij} + 2 + \rho_{ij}^{-1}$. Observe that v_{ij} is always upper bounded by $3 + \max(\rho_{ij}, \rho_{ji}) \le 3 + b \le 4b$, where we remind $b \ge \max_{i,j} \frac{w_i}{w_j}$.

We first argue that all F_{ij} are reasonably close to their expected values. For the sake of concision, we state the following assumptions about the constants, δ, k and the quantity $C_{n,\delta}$. Note that some of the intermediate results hold under weaker assumptions, but we omit these details for the sake of simplicity.

Assumption 1. *We have that* $\delta \le e^{-1}$, $C_{n,\delta} \ge c_1 \log(n/\delta)$, *and* $k \ge c_2 b(C_{n,\delta}+1)\max\{\Omega_{\max}, E_{\max}\}$.

The following lemma is a standard application of Chernoff's inequality. For completeness, a proof is included in Section **??** of the Supplementary Information.

Lemma 1. *There exist absolute constants constants* c_1, c_2 *such that, under Assumption 1, we have*

$$P\left(\max_{(i,j)\in E} |F_{ij} - p_{ij}| \ge \sqrt{\frac{C_{n,\delta}}{k v_{ij}}}\right) \le \delta.$$

The next lemma provides a convenient expression for the quantity $\log \hat{W} - \log w$ in terms of the "measurement errors" $F - p$. Note that the normalization assumption is not a loss of generality since w is defined up to a multiplicative constant, and is directly satisfied if $\hat{W}$ is obtained from (5).

Lemma 2. *Suppose w is normalized so that $\sum_{i=1}^n \log w_i = 0$. There exist absolute constants $c_1, c_2 > 0$ such that, under Assumption 1, there holds with probability $1-\delta$*

$$\log \hat{W} - \log w = L^\dagger BV(F-p) + L^\dagger B\Delta, \tag{7}$$

and

$$||\Delta||_\infty \leq O\left(\frac{bC_{n,\delta}}{k}\right), \tag{8}$$

where V is a $|E| \times |E|$ diagonal matrix whose entries are the v_{ij}, for all edges $(i,j) \in E$.

Proof. By definition

$$\log w_i - \log w_j = \log \rho_{ij} \text{ for all } (i,j) \in E,$$

which we can write as $B^T \log w = \log \rho$. It follows that

$$\log w = (BB^T)^\dagger B \log \rho = L^\dagger B \log \rho,$$

since w is assumed normalized so that $\sum_{i=1}^n \log w_i = 0$. Combining this with Eq. (5), we obtain

$$\log \hat{W} - \log w = L^\dagger B(\log R - \log \rho). \tag{9}$$

We thus turn our attention to analyzing the vector $\log R - \log \rho$. Our analysis will be conditioning on the event that for all $(i,j) \in E$,

$$\{|F_{ij} - p_{ij}| \leq \sqrt{\frac{C_{n,\delta}}{kv_{ij}}}\}, \tag{10}$$

which, by Lemma 1, holds with probability at least $1-\delta$. We will call this event $\mathcal{A}$.

We begin with one implication that comes from putting together event $\mathcal{A}$ and our assumption $k \geq c_1 bC_{n,\delta}$ (in Assumption 1) for a constant c_1 that we can choose: that we can assume that

$$\max_{(i,j)\in E} |F_{ij} - p_{ij}| \leq \frac{\min(p_{ij}, p_{ji})}{5}. \tag{11}$$

Indeed, from Eq. (10) for this last equation to hold it suffices to have $k \geq 25C_{n,\delta}/(v_{ij}p_{ij}^2)$ for all $(i,j) \in E$. Observing that

$$\frac{1}{v_{ij}p_{ij}^2} = \frac{1}{\frac{1}{p_{ij}p_{ji}}p_{ij}^2} = \rho_{ji} \leq b,$$

we see that assuming $k \geq 25bC_{n,\delta}$ is sufficient for Eq. (11) to hold conditional on event $\mathcal{A}$.

Our analysis of $\log R - \log \rho$ begins with the observation that since

$$R_{ij} = \frac{1-F_{ji}}{F_{ji}}, \;\; \rho_{ij} = \frac{1-p_{ji}}{p_{ji}}$$

we have that

$$\log R_{ij} - \log \rho_{ij} \;\; = \;\; \log\left(\frac{1}{F_{ji}} - 1\right) - \log\left(\frac{1}{p_{ji}} - 1\right)$$

Next we use Taylor's expansion of the function $g(x) = \log(1/x - 1)$, for which we have

$$g'(x) = \frac{1}{x(x-1)}, \;\; g'(p_{ji}) = -v_{ij}, \;\; g''(x) = \frac{1-2x}{x^2(1-x)^2}$$

to obtain that $\log R_{ij} - \log \rho_{ij}$ can thus be expressed as

$$-v_{ij}(F_{ji} - p_{ji}) + \frac{1}{2}\frac{1-2z_{ji}}{z_{ji}^2(1-z_{ji})^2}(F_{ij} - p_{ij})^2 \tag{12}$$

where z_{ji} lies between p_{ji} and F_{ji} (and $1 - z_{ji}$ lies thus between p_{ij} and F_{ij}). We can rewrite this equality in a condensed form

$$\log R - \log \rho = V(F-p) + \Delta, \tag{13}$$

where Δ corresponds to the second terms in (12), which we will now bound. Because we have conditioned on event $\mathcal{A}$, which, as discussed above implies $|F_{ji} - p_{ji}| \leq \min(p_{ji}, p_{ji})/5$, we actually have that $z_{ji} \in [0.8p_{ji}, 1.2p_{ji}]$ and that $1 - z_{ji}$ lying between p_{ij} and F_{ij} belongs to $[0.8p_{ij}, 1.2p_{ij}]$. Hence

$$|\Delta_{ij}| \leq \frac{1}{2}\frac{1}{0.8^4 p_{ij}^2 p_{ji}^2}(F_{ij} - p_{ij})^2 \leq c_3 v_{ij}\frac{C_{n,\delta}}{k},$$

for $c_3 = \frac{1}{2\times(0.8)^4}$, and where we have used (10) for the last inequality. Plugging this into Eq. (13) and (9) completes the proof, and Eq. (8) follows from the last equation combined with the fact that $v_{ij} \leq 4b$ for all $(i,j) \in E$. ■

The following lemma bounds how much the *ratios* of our estimates $\hat{W}_l$ differ from the corresponding ratios of the true weights w_l. To state it, we will use the notation

$$Q_{ij} = (\mathbf{e}_i - \mathbf{e}_j)(\mathbf{e}_i - \mathbf{e}_j)^T,$$

where $\mathbf{e}_i$ is the standard notation for the i'th basis vector. Furthermore, we define the product

$$\langle x, y\rangle_{(i,j)} = x^T B^T L^\dagger Q_{ij} L^\dagger By, \quad ||x||_{(i,j)}^2 = \langle x, x\rangle_{(i,j)}. \tag{14}$$

Observe that the matrix $B^T L^\dagger Q_{ij} L^\dagger B$ is positive semidefinite, which implies by standard arguments that

$$\langle x+y, x+y\rangle_{(i,j)} \leq 2\langle x, x\rangle_{(i,j)} + 2\langle y, y\rangle_{(i,j)}$$

holds for all vectors x, y.

Lemma 3. *Suppose w is normalized so that $\sum_{i=1}^n \log w_i = 0$. There exist absolute constants $c_1, c_2 > 0$ such, under Assumption 1, with with probability $1-\delta$, we have that for all pairs $i, j = 1, \ldots, n$,*

$$\left(\log \frac{\hat{W}_i}{\hat{W}_j} - \log \frac{w_i}{w_j}\right)^2 \leq 2\,||V(F-p)||_{(i,j)}^2 + 2\,||\Delta||_{(i,j)}^2\,, \tag{15}$$

and

$$||\Delta||_\infty \leq O\left(\frac{bC_{n,\delta}}{k}\right).$$

Proof. Observe that, on the one hand, using Lemma 2,

$$\begin{aligned}
&(\log \hat{W} - \log w)^T Q_{ij} (\log \hat{W} - \log w) \\
&= \left(L^\dagger BV(F-p) + L^\dagger B\Delta\right)^T Q_{ij} \left(L^\dagger BV(F-p) + L^\dagger B\Delta\right) \\
&= \langle V(F-p) + \Delta, V(F-p) + \Delta\rangle_{(i,j)} \\
&\leq 2\langle V(F-p), V(F-p)\rangle_{(i,j)} + 2\langle \Delta, \Delta\rangle_{(i,j)}
\end{aligned} \tag{16}$$

which is the right-hand side of (15). On the other hand, observe that

$$\begin{aligned}
&\left(\log \hat{W} - \log w\right) Q_{ij} \left(\log \hat{W} - \log w\right) \\
&= \left(\log \hat{W}_i - \log w_i - \left(\log \hat{W}_j - \log w_j\right)\right)^2 \\
&= \left(\log \tfrac{\hat{W}_i}{\hat{W}_j} - \log \tfrac{w_i}{w_j}\right)^2
\end{aligned} \tag{17}$$

Combining Eq. (16) with Eq. (17) completes the proof. ∎

Having proved Lemma 3, we now analyze each of the terms in the right-hand side of Eq. (15). We begin with the second term, i.e., with $||\Delta||^2_{(i,j)}$. To bound it, we will need the following inequality.

Lemma 4. *For any $\Delta \in \mathbb{R}^{|E|}$, we have that*

$$|\Delta^T B^T L^\dagger (\mathbf{e}_i - \mathbf{e}_j)| \leq ||\Delta||_\infty \sqrt{\Omega_{ij}|E_{ij}|},$$

where, recall, Ω_{ij} is the resistance between nodes i and j, and E_{ij} is the set of edges belonging to some simple path from i to j.

Proof. The result follows from circuit theory, and we sketch it out along with the relevant references. The key idea is that the vector $u = B^T L^\dagger (\mathbf{e}_i - \mathbf{e}_j)$ has a simple electric interpretation. We have that $u \in \mathbb{R}^{|E|}$ and the k'th entry of u is the current on edge k when a unit of current is put into node u at removed at node j. For details, see the discussion in Section 4.1 of (Vishnoi, 2013).

This lemma follows from several consequences of this interpretation. First, the entries of u are an acyclic flow from i to j; this follows, for example, from Thompson's principle which asserts that the current flow minimizes energy (see Theorem 4.8 of (Vishnoi, 2013)). Moreover, Thompson's principle further asserts that $\Omega_{ij} = ||u||_2^2$. Finally, by the flow decomposition theorem (Theorem 3.5 in (Ahuja et al., 2017)), we can decompose this flow along simple paths from i to j; this implies that $|\text{supp}(u)| \leq |E_{ij}|$.

With these facts in mind, we apply Cauchy-Schwarz to obtain

$$||u||_1 \leq ||u||_2 \sqrt{|\text{supp}(u)|} \leq \sqrt{\Omega_{ij}|E_{ij}|},$$

and then conclude the proof using Holder's inequality

$$|\Delta^T B^T L^\dagger (\mathbf{e}_i - \mathbf{e}_j)| = |\Delta^T u| \leq ||\Delta||_\infty ||u||_1 ||\Delta||_\infty \sqrt{\Omega_{ij}|E_{ij}|}.$$

∎

As a corollary, we are able to bound the second term in Eq. (15). The proof follows immediately by combining Lemma 4 with Lemma 3.

Corollary 1. *There exist absolute constants $c_1, c_2 > 0$ such that, under Assumption 1, with probability $1-\delta$, we have that for all pairs $i, j = 1, \ldots, n$,*

$$||\Delta||^2_{(i,j)} \leq O\left(\Omega_{ij} E_{ij} \frac{b^2 C_{n,\delta}^2}{k^2}\right).$$

We now turn to the first-term in Eq. (15), which is bounded in the next lemma.

Lemma 5. *There exist absolute constants c_1, c_2 such that, under Assumption 1, with probability $1-\delta$ we have that for all pairs $i, j = 1, \ldots, n$,*

$$||V(F-p)||^2_{(i,j)} \leq O\left(\Omega_{ij} \frac{b^2}{k}\left(1 + \log\frac{1}{\delta}\right)\right)$$

Proof. The random variable $X_{ij} - p_{ij}$ (where, recall, X_{ij} is the outcome of a single comparison between nodes i and j) is zero-mean and supported on an interval of length 1, and consequently it is subgaussian[4] with parameter 1 (see Section 5.3 of (Lattimore & Szepesvári, 2018)). By standard properties of subgaussian random variables, it follows that $v_{ij}(F_{ij} - p_{ij})$ is subgaussian with $\tau = v_{ij}/\sqrt{k} \leq 4b/\sqrt{k}$. It follows then from Theorem 2.1 of (Hsu et al., 2012) for subgaussian random variables applied to $||(e_i - e_j)B^T L^\dagger (F-p)||^2 = ||V(F-p)||^2_{(i,j)}$, that for any $t \geq 1$ there is a probability at least $1 - e^{-t}$ that

$$\begin{aligned}
||V(F-p)||^2_{(i,j)} &\leq \frac{16b^2}{k}\left(\text{tr}(M) + 2\sqrt{\text{tr}(M^2)t} + 2\,||M||\,t\right) \\
&\leq \frac{16b^2}{k}\text{tr}(M)(1+4t),
\end{aligned}$$

[4] A random variable Y is said to be subgaussian with parameter τ if $E[e^{\lambda Y}] \leq e^{\tau^2 \lambda^2/2}$ for all λ.

where we have used $\sqrt{t} \leq t$, $\mathrm{tr}(M^2) \leq \mathrm{tr}(M)^2$ and $||M|| \leq \mathrm{tr}(M)$. We now compute this trace.

$$\begin{aligned} \mathrm{tr}(M) &= \mathrm{tr}(B^T L^\dagger Q_{ij} L^\dagger B) \\ &= \mathrm{tr}(Q_{ij} L^\dagger B B^T L^\dagger) = \mathrm{tr}(Q_{ij} L^\dagger) \\ &= (\mathbf{e}_i - \mathbf{e}_j)^T L^\dagger (\mathbf{e}_i - \mathbf{e}_j) = \Omega_{ij}, \end{aligned} \tag{18}$$

where the second equality uses the well-known property of the Moore-Penrose pseudo-inverse: $A^\dagger A A^\dagger = A^\dagger$ for any matrix A (see Section 2.9 of (Drineas & Mahoney, 2018)); and last equality uses a well-known relation between resistances and Laplacian pseudoinverses, see Chapter 4 of (Vishnoi, 2013). The result follows then from the application of (18) to $t = \log 1/\delta$. ∎

Having obtained the bounds in the preceding sequence of lemmas, we now return to Lemma 3 and "plug in" the results we have obtained. The result is the following lemma.

Lemma 6. *There exist absolute constants $c_1, c_2 > 0$ such, under Assumption 1, with probability $1-\delta$, we have that for all pairs $i, j = 1, \ldots, n$,*

$$\left[\frac{\hat{W}_i}{\hat{W}_j} - \rho_{ij}\right]^2 \leq O\left(\rho_{ij}^2 \frac{b\Omega_{ij}}{k}\left(b(1+\log(1/\delta)) + \frac{bE_{i,j}C_{n,\delta}^2}{k}\right)\right)$$

Proof. By putting together Lemma 3 with Corollary 1 and Lemma 5, we obtain that, with probability at least $1-\delta$,

$$\begin{aligned} &\left(\log \frac{\hat{W}_i}{\hat{W}_j} - \log \frac{w_i}{w_j}\right)^2 \\ &\leq O\left(\frac{b\Omega_{ij}}{k}\left(b(1+\log(1/\delta)) + \frac{bE_{i,j}C_{n,\delta}^2}{k}\right)\right) \end{aligned} \tag{19}$$

Observe that for a sufficiently large c_2, if $k \geq c_2 E_{ij} C_{n,\delta}^2$ then the term $b(1+\log(1/\delta)) + \frac{bE_{i,j}C_{n,\delta}^2}{k}$ is bounded by $O(b(1+\log(1/\delta)))$. Hence, if k is also at least $c_2 b^2 \Omega_{ij}(1+\log(1/\delta))$ (which holds due to Assumption 1), equation (19) implies

$$\left|\log \frac{\hat{W}_i}{\hat{W}_j} - \log \frac{w_i}{w_j}\right| \leq 1. \tag{20}$$

A particular implication is that $\max\left(e^{\log(\hat{W}_i/\hat{W}_j)}, e^{\log(w_i/w_j)}\right) \leq e^{1+\log(w_i/w_j)}$. Applying the inequality $|e^a - e^b| \leq \max\{e^a, e^b\}|a-b|$ to (20) leads then to

$$\left|\frac{\hat{W}_i}{\hat{W}_j} - \frac{w_i}{w_j}\right| \leq e^{1+\log(w_i/w_j)}\left|\log \frac{\hat{W}_i}{\hat{W}_j} - \log \frac{w_i}{w_j}\right|$$

and now using $e^{\log(w_i/w_j)} = \rho_{ij}$, the proof follows by combining the last equation with Eq. (19). ∎

The next lemma demonstrates how to convert Lemma 6 into a bound on the relative error between $\hat{W}$ and the true weight vector w.

Lemma 7. *Suppose we have that*

$$\left[\frac{\hat{W}_i}{\hat{W}_j} - \rho_{ij}\right]^2 \leq \rho_{ij}^2 s_{ij}(k),$$

for all $i, j = 1, \ldots, n$. Fix index $\ell \in \{1, \ldots, n\}$. Then there hold

$$\sin(w, \hat{W}) \leq \max_j s_{j\ell}(k), \tag{21}$$

$$\sin(w, \hat{W}) \leq b^2 s_{\mathrm{avg}}, \tag{22}$$

where $s_{\mathrm{avg}} = \frac{\sum_{a,b=1,\ldots,n} s_{ab}}{n^2}$.

Proof. It follows from Lemma **??** that for all α,

$$\sin(w, \hat{W}) \leq \frac{||\hat{W} - \alpha w||_2^2}{||\alpha w||_2^2}.$$

Taking $\alpha = \hat{W}_\ell / w_\ell$, we get

$$\frac{||\hat{W} - \alpha w||_2^2}{||\alpha w||_2^2} = \frac{\sum_i (\hat{W}_i - \frac{\hat{W}_\ell}{w_\ell} w_i)^2}{\sum_i \frac{\hat{W}_\ell^2}{w_\ell^2} w_i^2} = \frac{\sum_i (\frac{\hat{W}_i}{\hat{W}_\ell} - \rho_{i\ell})^2}{\sum_i \rho_{i\ell}^2}.$$

Using the assumption of this lemma, we obtain

$$\frac{||\hat{W} - \alpha w||_2^2}{||\alpha w||_2^2} \leq \frac{\sum_i s_{i\ell}(k) \rho_{il}^2}{\sum_i \rho_{i\ell}^2}, \tag{23}$$

from which (21) follows. Another consequence of (23) is that

$$\frac{||\hat{W} - \alpha w||_2^2}{||\alpha w||_2^2} \leq \frac{(\max_i \rho_{i\ell}^2) \sum_i s_{i\ell}(k)}{n \min_j \rho_{j\ell}^2} \leq b^2 \frac{\sum_{i=1}^n s_{i\ell}}{n}, \tag{24}$$

where we used

$$\frac{\max_i \rho_{i\ell}}{\min_i \rho_{i\ell}} = \frac{\max_i w_i / w_\ell}{\min_j w_j / w_\ell} = \max_{i,j} \frac{w_i}{w_j} \leq b.$$

Observe now that since $s_{\mathrm{avg}} = \frac{1}{n} \sum_\ell \frac{1}{n} \sum_i s_{i\ell}$, there must exist at least one ℓ for which $\sum_{i=1}^n s_{i\ell} \leq s_{\mathrm{avg}}$. Hence (22) follows from (24). ∎

Having proven this last lemma, Theorem 1 follows immediately by combining By Lemma 6 and Lemma 7.

3. Experiments

The purpose of this section is two-fold. First, we would like to demonstrate that simulations are consistent with Theorem 1; in particular, we would like to see error scalings that are consistent with the average resistance, rather than e.g., spectral gap. Second, we wish observe that, although our results are asymptotic, in practice the scaling with resistance appears immediately, even for small k. Since our main contribution is theoretical, and since we do not claim that our algorithm is better than available methods in practice, we do not perform a comparison to other methods in the literature. Additional details about our experiments are provided in Section **??** in the Supplementary Information.

We begin with Erdos-Renyi comparison graphs. Figure 1 shows the evolution of the error with the number k of comparisons per edge. The error decreases as $O(1/\sqrt{k})$ as predicted. Moreover, this is already the case for small values of k.

Next we move to the influence of the graph properties. Figure 2 shows that the average error is asymptotically constant when n grows while keeping the expected degree $d := (n-1)p$ constant, and that it decreases as $O(1/\sqrt{d})$ when the expected degree grows while keeping n constant. This is consistent with our analysis in Table 1, and with the results (Boumal & Cheng, 2014) showing that the average resistance Ω_{avg} of Erdos-Renyi graphs evolves as $O(1/d)$.

We next consider lattice graphs in Figure 3. For the 3D lattice, the error appears to converge to a constant when n grows, which is consistent with our results since the average resistance of 3D lattice is bounded independently of n. The trend for the 2D lattice appears also consistent with a bound in $O(\sqrt{\log n})$ predicted by our results since the resistance on 2D lattice evolves as $O(\log n)$.

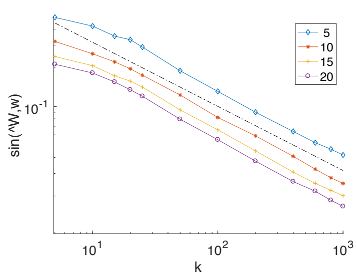

Figure 1. Error evolution with the number k of comparisons per edge in Erdos-Renyi graphs of 100 nodes, for different expected degrees $d = (n-1)p$, with $b = 10$. Each line corresponds to a different expected degree. The results are averaged over $N_{\text{test}} = 100$ tests. The dashed line is proportional to $1/\sqrt{k}$.

4. Conclusion

Our main contribution has been to demonstrate, by a combination of upper and lower bounds, that the error in quality estimation from pairwise comparisons scales as the graph resistance. Our work motivates a number of open questions.

First, our upper and lower bounds are not tight with respect to skewness measure b. We conjecture that the scaling of $\widetilde{O}(\sqrt{b\Omega_{\text{avg}}/k})$ for relative error is optimal, but either

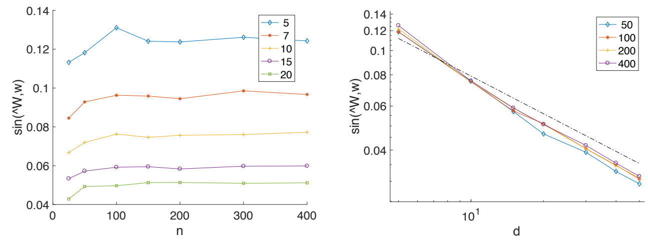

(a) (b)

Figure 2. Error evolution with the number of nodes n for different expected degrees $d = (n-1)p$ (a), and with the expected degree $(n-1)p$ for different number of nodes n (b). There are $k = 100$ comparisons per edge, $b = 5$, and results are averaged over $N_{\text{test}} = 50$ tests. The dashed line in (b) is proportional to $1/\sqrt{k}$.

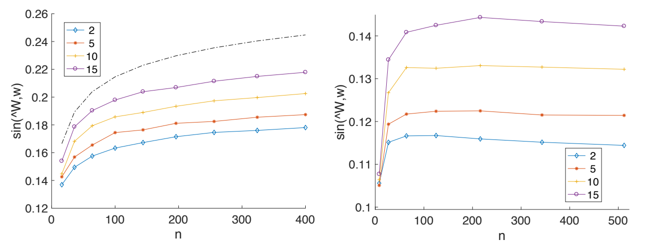

(a) (b)

Figure 3. Error evolution with the number of nodes for regular lattices in 2D (a), and 3D (b). Each line corresponds to a different choice of b. The number of comparisons is $k = 100$ per edge. Results are averaged over respectively $N_{\text{test}} = 1000$ and $N_{\text{test}} = 2000$ tests. The dashed line in (a) is proportional to $\sqrt{\log n}$.

upper of lower bounds matching this quantity are currently unknown.

Second, it would interesting to obtain non-asymptotic version of the results presented here. Our simulations are consistent with the asymptotic scaling $\widetilde{O}(\sqrt{\Omega_{\text{avg}}/k})$ (ignoring the dependence on b) being effective immediately, but at the moment we can only prove this scaling governs the behavior as $k \to \infty$.

References

Agarwal, A., Agarwal, S., Assadi, S., and Khanna, S. Learning with limited rounds of adaptivity: Coin tossing, multi-armed bandits, and ranking from pairwise comparisons. In *Conference on Learning Theory*, pp. 39–75, 2017.

Ahuja, R. K., Magnanti, T. L., and Orlin, J. B. *Network Flows: Theory, Algorithms, and Applications*. Pearson Education, 2017.

Ajtai, M., Feldman, V., Hassidim, A., and Nelson, J. Sorting and selection with imprecise comparisons. *ACM Transactions on Algorithms (TALG)*, 12(2):19, 2016.

Avin, C. and Ercal, G. On the cover time and mixing time of random geometric graphs. *Theoretical Computer Science*, 380(1):2, 2007.

Boumal, N. and Cheng, X. Concentration of the Kirchhoff index for Erdős–Rényi graphs. *Systems & Control Letters*, 74:74–80, 2014.

Bradley, R. A. and Terry, M. E. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

Braverman, M. and Mossel, E. Sorting from noisy information. *arXiv preprint arXiv:0910.1191*, 2009.

Chen, Y. and Suh, C. Spectral mle: Top-k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*, pp. 371–380, 2015.

Drineas, P. and Mahoney, M. W. Lectures on Randomized Numerical Linear Algebra. *The Mathematics of Data*, 25: 1, 2018.

Duchi, J. Assouads method. https://web.stanford.edu/class/stats311/Lectures/lec-04.pdf. Lecture notes.

Falahatgar, M., Orlitsky, A., Pichapati, V., and Suresh, A. T. Maximum selection and ranking under noisy comparisons. *arXiv preprint arXiv:1705.05366*, 2017.

Hajek, B. and Raginsky, M. Statistical learning theory. http://maxim.ece.illinois.edu/teaching/SLT/SLT.pdf. Book draft.

Hajek, B., Oh, S., and Xu, J. Minimax-optimal inference from partial rankings. In *Advances in Neural Information Processing Systems*, pp. 1475–1483, 2014.

Hsu, D., Kakade, S., and Zhang, T. A Tail Inequality for Quadratic Forms of subgaussian Random Vectors. *Electronic Communications in Probability*, 17, 2012.

Jannach, D., Resnick, P., Tuzhilin, A., and Zanker, M. Recommender systems – beyond matrix completion. *Communications of the ACM*, 59(11):94–102, 2016.

Jiang, X., Lim, L.-H., Yao, Y., and Ye, Y. Statistical ranking and combinatorial hodge theory. *Mathematical Programming*, 127(1):203–244, 2011.

Landau, H. and Odlyzko, A. Bounds for eigenvalues of certain stochastic matrices. *Linear algebra and its Applications*, 38:5–15, 1981.

Lattimore, T. and Szepesvári, C. Bandit Algorithms. http://downloads.tor-lattimore.com/banditbook/book.pdf, 2018. Book draft.

Luce, R. D. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.

Mallows, C. L. Non-null ranking models. i. *Biometrika*, 44 (1/2):114–130, 1957.

Mao, C., Weed, J., and Rigollet, P. Minimax rates and efficient algorithms for noisy sorting. *arXiv preprint arXiv:1710.10388*, 2017.

Maystre, L. and Grossglauser, M. Just sort it! A simple and effective approach to active preference learning. *arXiv preprint arXiv:1502.05556*, 2015.

Negahban, S., Oh, S., and Shah, D. Iterative ranking from pair-wise comparisons. In *Advances in neural information processing systems*, pp. 2474–2482, 2012.

Negahban, S., Oh, S., and Shah, D. Rank centrality: Ranking from pairwise comparisons. *Operations Research*, 65 (1):266–287, 2016.

Pananjady, A., Mao, C., Muthukumar, V., Wainwright, M. J., and Courtade, T. A. Worst-case vs average-case design for estimation from fixed pairwise comparisons. *arXiv preprint arXiv:1707.06217*, 2017.

Rajkumar, A. and Agarwal, S. A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *International Conference on Machine Learning*, pp. 118–126, 2014.

Shah, N. B., Balakrishnan, S., Bradley, J., Parekh, A., Ramchandran, K., and Wainwright, M. J. Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *The Journal of Machine Learning Research*, 17(1):2049–2095, 2016.

Shah, N. B., Balakrishnan, S., Guntuboyina, A., and Wainwright, M. J. Stochastically transitive models for pairwise comparisons: Statistical and computational issues. *IEEE Transactions on Information Theory*, 63(2):934–959, 2017.

Spielman, D. A. and Teng, S.-H. Nearly linear time algorithms for preconditioning and solving symmetric, diagonally dominant linear systems. *SIAM Journal on Matrix Analysis and Applications*, 35(3):835–885, 2014.

Sylvester, J. A. Random walk hitting times and effective resistance in sparsely connected erdos-renyi random graphs. *arXiv preprint arXiv:1612.00731*, 2016.

Szörényi, B., Busa-Fekete, R., Paul, A., and Hüllermeier, E. Online rank elicitation for plackett-luce: A dueling bandits approach. In *Advances in Neural Information Processing Systems*, pp. 604–612, 2015.

Tao, T. *Topics in random matrix theory*, volume 132. American Mathematical Soc., 2012.

Vishnoi, N. $Lx = b$. *Foundations and Trends in Theoretical Computer Science*, 8, 2013.

Wilf, H. S. The editor's corner: the white screen problem. *The American Mathematical Monthly*, 96(8):704–707, 1989.

Yue, Y., Broder, J., Kleinberg, R., and Joachims, T. The k-armed dueling bandits problem. *Journal of Computer and System Sciences*, 78(5):1538–1556, 2012.