
Generalized Results for the Existence and Consistency of the MLE in the Bradley-Terry-Luce Model

Heejong Bong¹ Alessandro Rinaldo¹

Abstract

Ranking tasks based on pairwise comparisons, such as those arising in online gaming, often involve a large pool of items to order. In these situations, the gap in performance between any two items can be significant, and the smallest and largest winning probabilities can be very close to zero or one. Furthermore, each item may be compared only to a subset of all the items, so that not all pairwise comparisons are observed. In this paper, we study the performance of the Bradley-Terry-Luce model for ranking from pairwise comparison data under more realistic settings than those considered in the literature so far. In particular, we allow for near-degenerate winning probabilities and arbitrary comparison designs. We obtain novel results about the existence of the maximum likelihood estimator (MLE) and the corresponding ℓ_2 estimation error without the bounded winning probability assumption commonly used in the literature and for arbitrary comparison graph topologies. Central to our approach is the reliance on the Fisher information matrix to express the dependence on the graph topology and the impact of the values of the winning probabilities on the estimation risk and on the conditions for the existence of the MLE. Our bounds recover existing results as special cases but are more broadly applicable.

1. Introduction

The task of eliciting a ranking over a possibly large collection of items based on partially observed pairwise comparisons is of great relevance for an increasingly broader range

of applications, including sports analytics, multi-player online games, collection/analysis of crowdsourced data, marketing, aggregation of online search engine response queries and social choice models, to name a few. In response to these growing needs for theories and methods for high-dimensional ranking problems, the statistical and machine learning literature has witnessed a recent flurry of activities and novel results concerning high-dimensional parametric ranking models. Among them, the renown Bradley-Terry-Luce (BTL) model (Bradley & Terry, 1952; Luce, 2012) has experienced a surge in popularity and use, due to its high interpretability and expressive power.

Though the statistical literature on the BLT model is extensive, dating back to the seminal work on existence of the MLE by (Ford Jr, 1957), a series of recent contributions have led to accurate upper and lower bounds on the estimation error of the BTL model parameters under various high-dimensional settings. See Related Work below. These bounds, as well as most other theoretical results in the literature, are typically predicated on the key assumption that the pairwise winning probabilities remain bounded away from 0 and 1 by a fixed and even known amount. However, in high-dimensional ranking modeling, it is desirable to allow for the winning probabilities to become degenerate (i.e. to approach 0 or 1) as the number of items to be compared grows. This feature expresses characteristics found in real-life ranking tasks, such as ratings in massive multi-player online games. In those cases, it is expected that the maximal performance gap – also known as dynamic range – among different players increases with the number of players, so that the best players will be vastly superior to those with poor skills. Another significant complication frequently arising in real-life ranking tasks is that, due to experimental or financial constraints, only few among all possible pairwise comparisons are observed, resulting in a possibly irregular or sparse comparison graph. The combined effect of a broad gap in performance and of a possibly unfavorable comparison graph topology may have a considerable impact on the estimability of the BTL model parameters. The main goal of this paper is to quantify such impact in novel and explicit ways. We will expand on some of the latest theoretical contributions to parametric ranking using the BTL model in order to (i) formulate novel conditions for the existence of

¹Department of Statistics and Data Sciences, Carnegie Mellon University, Pittsburgh, PA, USA. Correspondence to: Heejong Bong <hbong@andrew.cmu.edu>, Alessandro Rinaldo <arinaldo@cmu.edu>.

the MLE of the model parameters and (ii) derive more general and sharper ℓ_2 estimation bounds. In particular, we will relax the prevailing assumption of non-degenerate winning probabilities frequently used in the literature on ranking while allowing for arbitrary comparison graph topologies. Our bounds hold under much broader settings than those considered so far and thus offers theoretical validation for the use of the BTL model in real-life problems.

This paper makes the following two main contributions.

- In Section 3, we present a novel sufficient condition for the existence of the MLE of the parameters of the BTL model under an arbitrary comparison graph and dynamic ranges that encompass and greatly generalize previous results. Existence of the MLE is a necessary pre-requisite for estimability of the model parameters, and therefore it is important to identify these minimal conditions for which the ranking task is at least well-posed. Compared to the few existing results in the literature, which hold only under very specific settings, our condition accounts for the effect of the graph topology and of the performance gap and, despite its generality, immediately recovers the few known results as special cases; see Simons & Yao (1999), Yan et al. (2012), and Han et al. (2020).
- In Section 4, we present a new ℓ_2 consistency rate for the MLE without bounded dynamic range and again permitting arbitrary graph topologies that generalizes the ℓ_2 error bound of (Shah et al., 2016) and is more generally applicable. Two key technical contributions play a central role in our analysis: (i) the use of a proxy loss function that, unlike the negative log-likelihood loss, will retain strong convexity behavior even if the dynamic range is not bounded and (ii) the reliance of the Fisher information to adequately capture the efficiency of the MLE and its dependence on both the graph topology and the performance gap. In Section 4.1 we illustrated in detail how our new bound improves the current state-of-the-art result of (Shah et al., 2016) and is tighter in more general and challenging ranking settings. These findings are further corroborated in simulation studies.

Recent Related Work. Over the last few years the high-dimensional statistics literature has produced an impressive body of new results on the BTL model, demonstrating its performance and effectiveness in a variety of complex ranking tasks scenarios. We present a non-exhaustive list of the related works. Novel bounds on the ℓ_2 estimation of the BLT models are derived in Negahban et al. (2012); Hajek et al. (2014); Khetan & Oh (2016); Shah et al. (2016); Chen et al. (2020a). Hendrickx et al. (2019; 2020) consider instead the sine error measure, while Chen et al. (2019; 2020a;b) focus

on the ℓ_∞ estimation error. To date, the sharpest ℓ_2 related estimation rates are those obtained by Shah et al. (2016); Hendrickx et al. (2020), and exhibit an explicit dependence on the topology of the comparison graph, i.e. the undirected graph over the items to be ranked, with edges indicating the pairs of items that are compared.

Notation. For symmetric matrices $A, B \in \mathbb{R}^{d \times d}$, we write $A \succ B$ ($A \succeq B$) if $A - B$ is positive (semi-)definite. Also, for a vector $x \in \mathbb{R}^d$, $\|x\|_A^2 := x^\top A x$. We denote the eigenvalues of a positive semi-definite matrix A by $0 \leq \lambda_1(A) \leq \lambda_2(A) \leq \dots \leq \lambda_d(A)$. We will use repeatedly the well known fact that if A is the graph Laplacian or the Fisher information matrix induced by the BTL model then $A\mathbf{1} = \mathbf{0}$ and hence $\lambda_1(A) = 0$. Finally, for positive real sequences x_d and y_d with respect to d , we write $x_d = o(y_d)$ if $\frac{x_d}{y_d} \rightarrow 0$, $O(y_d)$ if $\frac{x_d}{y_d}$ is upper bounded, $\omega(y_d)$ if $\frac{x_d}{y_d} \rightarrow \infty$, and $\Omega(y_d)$ if $\frac{x_d}{y_d}$ is lower bounded as $d \rightarrow \infty$. In particular, we say $x_d = \Theta(y_d)$ if $x_d = O(y_d)$ and $x_d = \Omega(y_d)$. Accordingly, we write $x_d = o_p(y_d), O_p(y_d), \omega_p(y_d)$, or $\Omega_p(y_d)$ if the respective condition meets in probability; e.g., $x_d = O_p(y_d)$ if $\mathbb{P}[\frac{x_d}{y_d} \leq c] \rightarrow 0$ as $d \rightarrow \infty$ for some constant $c > 0$.

2. Bradley-Terry-Luce Model for Pairwise Comparisons

In this section, we review the Bradley-Terry-Luce model for ranking with pairwise comparisons introduced by Bradley & Terry (1952) and Luce (2012). In this model, the d items to be ranked are each assigned an unknown numeric *quality scores*, denoted by the vector $w^* = (w_1^*, \dots, w_d^*) \in \mathbb{R}^d$ so that the probability that item i is preferred to item j in a pairwise comparison is specified as

$$\mathbb{P}[i \text{ defeats } j] = \frac{e^{w_i^*}}{e^{w_i^*} + e^{w_j^*}}. \quad (1)$$

To ensure identifiability of the model, we impose the additional constraint that $\sum_{i=1}^d w_i^* = 0$. The difference in performance between the best and worst items, that is, $\max_{i,j} |w_i^* - w_j^*|$, is referred to as the maximal performance gap or *dynamic range*.

Data are collected in the form of n pairwise comparisons among the items, which are modeled as a sequence of independent labeled Bernoulli random variables (y_1, \dots, y_n) , where the label for y_k is the pair of distinct items $(i_k, j_k) \in [d] \times [d]$ that have been compared. Accordingly, for $k = 1, \dots, n$,

$$y_k = \begin{cases} 1 & \text{w.p. } \frac{\exp(w_{i_k}^*)}{\exp(w_{i_k}^*) + \exp(w_{j_k}^*)}, \\ -1 & \text{w.p. } \frac{\exp(w_{j_k}^*)}{\exp(w_{i_k}^*) + \exp(w_{j_k}^*)}. \end{cases} \quad (2)$$

Then, the random variable $X_k = y_k(\mathbf{e}_{i_k} - \mathbf{e}_{j_k}) \in \mathbb{R}^d$ indicates the result of the k -th comparison between items i_k and j_k : the item with entry 1 won while the other with entry -1 was defeated. We record which items have been compared and the number of comparisons between any two items using the comparison matrix X , whose k -th row is X_k^\top . Then, the matrix

$$L := \frac{1}{n} X^\top X. \quad (3)$$

is the normalized Laplacian of the *comparison graph*, the weighted undirected graph over the items in which the weight of the (i, j) edge is the fraction of times, out of the n comparisons, that items i and j were compared. We note that L turns out to be a function of i_k 's and j_k 's – i.e., $L = \frac{1}{n} \sum_{k=1}^n (e_{i_k} - e_{j_k})(e_{i_k} - e_{j_k})^\top$ – so that L is deterministic and also invariant with respect to the order between i_k and j_k .

We estimate the quality scores and then rank the items based on the maximum likelihood estimator. The (normalized) log-likelihood $l(w)$ of the BTL model is written in terms of the X_k 's as

$$l(w) = -\frac{1}{n} \sum_{k=1}^n \log(1 + e^{-\langle w, X_k \rangle}), \quad (4)$$

and the maximum likelihood estimator (MLE) \hat{w} is obtained by solving the optimization problem

$$\hat{w} = \arg \min_{w \in \mathbb{R}^d} -l(w) \quad \text{s.t.} \quad \sum_{i=1}^d w_i = 0. \quad (5)$$

Notice that the the supremum in the above expression may not be attained, in which case we say the MLE does not exist.

To obtain a tight consistency rate for the MLE, we need to take into account the both information about the topology of the comparison graph and the distribution of performance across items. As our results demonstrate, the Fisher information matrix efficiently captures the impact of both sources of statistical difficulty. We note that under the BTL model, the Hessian matrix of the likelihood is constant with respect to X and the Fisher information is easily calculated as

$$\begin{aligned} \mathcal{I}_{ij}^* &:= [\mathcal{I}(w^*)]_{ij} = \mathbb{E}[-\nabla^2 l(w^*)]_{ij} \\ &= \begin{cases} \frac{L_{ij}}{(e^{(w_i^* - w_j^*)/2} + e^{(w_j^* - w_i^*)/2})^2}, & i \neq j, \\ -\sum_{k:k \neq i} [\mathcal{I}(w^*)]_{ik}, & i = j. \end{cases} \end{aligned} \quad (6)$$

One may easily recognize that the (i, j) -th entry of $\mathcal{I}(w^*)$ is equal to the corresponding entry of the normalized Laplacian L weighted by $(e^{(w_i^* - w_j^*)/2} + e^{(w_j^* - w_i^*)/2})^{-2}$, which is decreasing in the performance gap $|w_i^* - w_j^*|$. In the

following sections, we will derive a novel; sufficient condition for the MLE existence and a bound on the magnitude of its error using the Fisher information \mathcal{I}^* instead of the Laplacian L .

3. Existence of the MLE under an Arbitrary Graph Topology

The MLE defined in Eq. (5) often plays a fundamental role in any ranking task based on the BTL model. However, such an optimization-based estimator is not generally guaranteed to exist. For example, suppose that there is a partition of the items into two or more group such that there is no comparison across different groups *by design*. Equivalently, the comparison graph is *disconnected*. In this case, no amount of data will be able to resolve the ranking between any two items belonging to different groups, and the model is non-identifiable. As a result, the MLE does not exist. To avoid such cases, we will assume throughout that the comparison graph is fully connected, as is commonly done in the literature. Under this assumption, the Laplacian and Fisher information matrices have only one zero eigenvalue (with associated eigenvector spanning the one-dimensional subspace comprised of vectors with constant entries), and hence $\lambda_2(L)$ – known as the *algebraic connectivity* of the graph (Brouwer & Haemers, 2011) – and $\lambda_2(\mathcal{I}^*)$ are always positive.

Existence and uniqueness of the solution of Eq. (5) is still not guaranteed even with a connected comparison graph and demand stronger conditions. Indeed, when one item wins all comparisons against all the others, the log-likelihood function will not admit a supremum over \mathbb{R}^n as the optimum will be achieved when a positive infinity score is assigned to the undefeated item. To avoid such cases, no single item nor group of items should overwhelm the others. As proved by Zermelo (1929) and Ford Jr (1957), this is indeed sufficient and necessary for a unique solution. We formalize this condition below.

Condition 3.1 (Ford Jr (1957)). In every possible partition of the items into two nonempty subsets, some item in the first set has beaten another in the second set. Or equivalently, for each ordered pair (i, j) , there exists a sequence of indices $i_0 = i, i_1, \dots, i_m = j$ such that i_{k-1} have beaten i_k for $k = 1, \dots, m$.

Condition 3.1 is informative about the existence of the MLE only after the results of all the comparisons are observed. An important yet not fully established problem is the prediction of Condition 3.1 ahead of observing any data. That is, we seek conditions for the score w^* and the comparison graph, encoded by the Laplacian L , to generate comparison results satisfying Condition 3.1 with high probability. To the best of our knowledge, this problem has been studied in only few, limited settings and a general analytic result is lacking.

Simons & Yao (1999) stated and proved a sufficient condition for the complete comparison graph, which demands a sufficient amount of comparisons and small dynamic range. This condition was further extended by (Yan et al., 2012) to allow for sparse graph topologies. More recently, (Han et al., 2020) significantly sharpened Simon’s and Yao’s analysis and provided a sufficient condition for sparse Erdős-Rényi type of graph topologies (Erdős & Rényi, 1960). In our first result, we further refine the arguments of Simons & Yao (1999) to derive a novel and general sufficient condition for existence of the MLE that is valid for arbitrary comparison graphs and depends on the smallest non-zero eigenvalue of the Fisher information matrix. See Appendix A.1 for the proof.

Theorem 3.2. *If $\lambda_2(\mathcal{I}^*) \geq 2\frac{\log d}{n}$, then*

$$\mathbb{P}[\text{Condition 3.1 fails}] \leq \frac{2}{\sqrt{d}}. \quad (7)$$

We emphasize that the above sufficient condition is not only generally applicable to any graph topology but in fact, as we will show next, recovers existing results as special cases. We believe that the strength and generality of our result stem from the direct use of the Fisher information matrix.

3.1. Comparison to Existing Results

The complete graph. Under complete comparison graphs, Theorem 3.2 recovers Lemma 1 of Simons & Yao (1999). In the setting, every pair of the items are compared the same number of times. Let $M := \max_{i,j \in [d]} \exp(w_i^* - w_j^*) - 1$ a quantity analogue to the term M_t in Simons & Yao (1999) – represent the dynamic range of the compared items. Then, according to Eq. (6), every off-diagonal element in \mathcal{I}^* has absolute value larger than $\frac{M}{(1+M)^2}$, and it is easy to show that $\mathcal{I}^* \succeq \frac{M}{(1+M)^2} L$. Since the algebraic connectivity of the complete graph is given by $\lambda_2(L) = \frac{2}{d-1}$ (Brouwer & Haemers, 2011), we arrive at the bound

$$\lambda_2(\mathcal{I}^*) \geq \frac{M}{(1+M)^2} \lambda_2(L) \geq \frac{1}{1+M} \frac{1}{d-1}. \quad (8)$$

As a result, the sufficient condition $M = o\left(\frac{d}{\log d}\right)$ of Lemma 1 of (Simons & Yao, 1999) implies that $\lambda_2(\mathcal{I}^*) = \omega\left(\frac{\log d}{n}\right)$ which by Theorem 3.2 implies existence of the MLE (to be precise, Lemma 1 of (Simons & Yao, 1999) assumes the more restrictive scaling $M = o\left(\sqrt{\frac{d}{\log d}}\right)$, but in fact their result holds under the weaker scaling reported above).

The Erdős Rényi graph. As a second application, assume that the comparison graph is an Erdős-Rényi graph with parameter p , possibly vanishing in d , a setting that has been the focus of much of the recent work on the BTL model

(Negahban et al., 2012; Yan et al., 2012; Han et al., 2020). In this case, the sharpest known sufficient condition for the existence of the MLE of the BTL model parameters is due to (Han et al., 2020), and requires that $M = o\left(\frac{dp}{\log d}\right)$, provided that $p = \omega\left(\frac{\log d}{d}\right)$ and assuming again for simplicity a constant number of comparisons for each pair; see also (Yan et al., 2012) for an earlier result requiring p to be bounded away from 0. Under the slightly weaker requirement that p is of order $\Omega\left(\frac{\log d}{d}\right)$, the bound in (8) along with the fact that $\lambda_2(L) = \Theta_p\left(\frac{dp}{n}\right)$ (a fact that can be deduced from Kolokolnikov et al., 2014, after noting that we use a different normalization for the Laplacian) shows that the same scaling of $M = o\left(\frac{dp}{\log d}\right)$ found by Han et al. (2020) does in fact satisfies the condition of Theorem 3.2.

4. Consistency of MLE without the Bounded Dynamic Range Assumption

To the best of our knowledge, the work by Shah et al. (2016) was the first attempt to study consistency of the MLE of the BTL parameters in the ℓ_2 and weighted ℓ_2 loss under arbitrary comparison graphs. Their rates are complemented by minimax lower bounds that also exhibit an explicit dependence on comparison graph topology, and can be used to identify comparison graph topologies that produce minimax optimal estimation rates. These results crucially require the assumption, frequently imposed in the literature on ranking, of a *bounded dynamic range*, namely that $\|w^*\|_\infty < B$, for a *known* value B . Under these constraint, the feasible space for the likelihood optimization problem Eq. (5) is a compact set $\{x \in \mathbb{R}^d : \|x\|_\infty \leq B\}$, over which the negative log-likelihood is by construction strongly convex. In particular, the resulting regularized MLE is guaranteed to exist, even in cases in which the MLE may not. Next, using a Taylor series expansion, the authors obtain at a strongly quadratic approximation of the log-likelihood in a neighborhood of the true parameter value, from which the error upper bounds of their Theorems 1 and 2 follow. However, the assumption imposed by Shah et al. (2016) is fairly restrictive and unrealistic in high-dimensional settings. Furthermore, the use of the constrained MLE is only feasible when prior information about the value of the hyper-parameter B is available, a condition that is unrealistic in many settings. Hence, a consistency rate for the MLE that does not depend on the bounded dynamic range assumption is in order.

Without the aid of the bounded dynamic range assumption, the summand $\log(1 + e^{-(w^*, X_k)})$ in the negative log-likelihood function $w \in \mathbb{R}^d \mapsto -l(w)$ gets increasingly flatter as the absolute value of the exponent increases, and hence the curvature of $-l(\cdot)$ is no longer uniformly lower-bounded. In fact, the restricted curvature assumption is central to the analysis of Shah et al. (2016) and their tech-

niques no longer apply when $-l(\cdot)$ is not strongly convex. We overcome this hurdle by focusing instead on a proxy function $h_d : \mathbb{R}^d \rightarrow \mathbb{R}^d$ for the negative log-likelihood that offers two key advantages: (i) it turns $-l(\cdot)$ into a strongly convex loss function with respect to $h(\cdot - w^*)$ and (ii) it relates to the Fisher information matrix through the inequality

$$\mathbb{E}[l(w^*) - l(w)] \geq h_d(w - w^*)^\top \mathcal{I}^* h_d(w - w^*). \quad (9)$$

With these critical modifications in place, the ℓ_2 consistency rate of Theorem 4.1, the main result of this section, follows using similar proof arguments as those deployed in Shah et al. (2016).

Before stating our main result, we first define and study the proxy function h_d . First, let $h : \mathbb{R} \rightarrow \mathbb{R}$ be the univariate real function defined by $h(x) := \text{sgn}(x)(\sqrt{|x| + 1} - 1)$. See Fig. 1a. We point out that h is an increasing function that behaves linearly near the origin and whose slope decreases at the same rate as the slope of $\sqrt{|x|}$ for $x \rightarrow \pm\infty$. Furthermore, it can be verified that the function $x \mapsto h(x)^2$ is upper-bounded and lower-bounded by $|x|$ and $\frac{|x|}{2} - 1$, respectively; see Fig. 1b. As a result, $h(x)^2$ behaves like a quadratic function of x near zero while increasing linearly for sufficiently large $|x|$. It is precisely this property that ensures strong convexity of $l(w)$ with respect to the proxy function $h_d(w - w^*)$, which we define next. We take h_d to be the d -dimensional element-wise evaluation of h . That is, $[h_d(x)]_i = h(x_i)$, $i = 1, \dots, d$. As remarked above, when $x \rightarrow \infty$, $\|h_d(x)\|_2^2 = \sum_{i=1}^d h(x_i)^2$ behaves like a linear function, so its curvature converges to zero. Hence, although $w \in \mathbb{R}^d \mapsto -l(w)$ is not itself strongly convex, the MLE problem in Eq. (5) can be approximated by a quadratic problem with respect to $w \in \mathbb{R}^d \mapsto h_d(w - w^*)$.

The ℓ_2 convergence rate of the MLE derived below depends on a parameter κ defined by

$$\kappa := \lambda_d(L^{\frac{1}{2}} \mathcal{I}^{*+} L^{\frac{1}{2}}) \quad (10)$$

where \mathcal{I}^{*+} is the Moore-Penrose pseudo-inverse of the Fisher information matrix. As a function of \mathcal{I}^* , κ takes into account both the comparison graph topology and the gaps in performance among the items. Importantly, κ is determined only by the performance gaps among compared items, unlike the the dynamic range, which is the maximal performance gap across *all* items, including those that are never compared. We will elucidate the gains from replacing the traditional dynamic range parameter by the new parameter κ through a concrete example below, when we compare our bounds to existing results.

Theorem 4.1. *Assume that $\lambda_2(\mathcal{I}^*) \geq c_0 \frac{\sqrt{v_{\max} \log d}}{n}$ for some universal constant $c_0 > 0$, where v_{\max} be the maximum degree of the comparison graph. Then, the MLE \hat{w} for the BTL model exists with probability at least $1 - \frac{1}{d}$, and the*

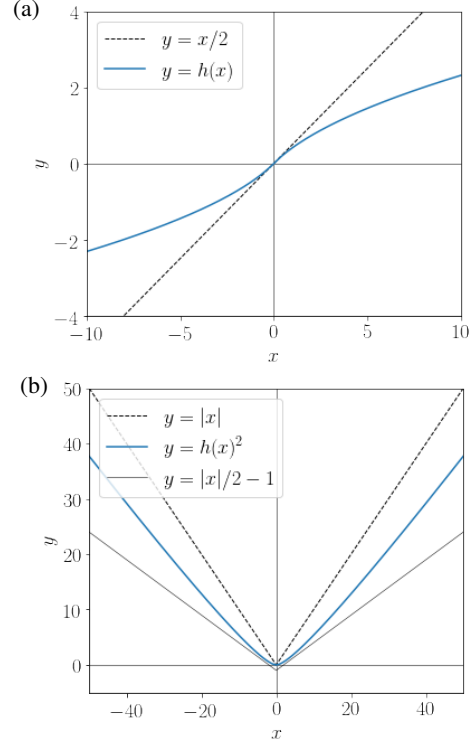


Figure 1. Shape of the proxy function, h , and its square. (a) The absolute value of $h(x)$ is smaller than $\frac{|x|}{2}$, and hence $h(x)^2$ is upper bounded by $\frac{x^2}{4}$. (b) Also, $h(x)^2$ is upper- and lower bounded by $|x|$ and $\frac{|x|}{2} - 1$, respectively. In sum, $h(x)^2$ behaves like a quadratic function around zero and an absolute function for sufficiently large $|x|$.

error $\Delta = \hat{w} - w^*$ satisfies

$$\|h_d(\Delta)\|_{\mathcal{I}^*}^2 \leq c_1 \kappa \frac{td}{n}, \quad (11)$$

and, as a corollary,

$$\|h_d(\Delta)\|_2^2 \leq c_1 \frac{\kappa}{\lambda_2(\mathcal{I}^*)} \frac{td}{n}, \quad (12)$$

with probability at least $1 - e^{-t} - \frac{1}{d}$ for any $t > 0$ and some universal constant $c_1 > 0$.

The only assumption in the theorem is the requirement that $\lambda_2(\mathcal{I}^*) \geq c_0 \frac{\sqrt{v_{\max} \log d}}{n}$, which is imposed to ensure existence of the MLE and a sufficient degree of curvature of the log-likelihood. As remarked above, this issue does not arise with the regularized MLE, which is always well defined.

We point out that the condition $\lambda_2(\mathcal{I}^*) \geq c_0 \frac{\sqrt{v_{\max} \log d}}{n}$ is in general stronger than the sufficient condition for the existence of the MLE implied by Theorem 3.2. To see this, we use the well known fact that the maximal degree is of

the same order as the largest eigenvalue of the graph Laplacian, which in our settings implies that $v_{\max} = \Theta(\lambda_2(\mathcal{I}^*)n)$. Thus, $\lambda_2(\mathcal{I}^*) \geq c_0 \frac{\sqrt{v_{\max} \log d}}{n}$ is equivalent to $\lambda_2(\mathcal{I}^*) = \Omega\left(\frac{\lambda_d(\mathcal{I}^*)}{\lambda_2(\mathcal{I}^*)} \frac{\log d}{n}\right)$. Ignoring constants, the latter condition is stricter than the condition $\lambda_2(\mathcal{I}^*) \geq 2 \frac{\log d}{n}$ from Theorem 3.2, unless all the non-zero eigenvalues of \mathcal{I}^* are of the same order. It is an open problem to fill the gap between the conditions for the MLE existence and ℓ_2 error rate guarantee.

Below we provide detailed comparisons to existing results and explain how Theorem 4.1 delivers significant improvements by allowing for a significant weakening of the assumptions.

4.1. Comparison to Existing Results

Erdős-Rényi graph under bounded dynamic range assumption. We first compare Theorem 4.1 to Theorem 3.1 of Chen et al. (2019), who provide an ℓ_2 rate for the BTL parameters assuming an Erdős-Rényi comparison graph with edge probability p and bounded dynamic range B . Theorem 3.1 of Chen et al. (2019) requires that $p = \Omega\left(\frac{\log d}{d}\right)$, to ensure that the comparison graph is connected. This additional constraint is equivalent to the condition $\lambda_2(\mathcal{I}^*) \geq c_0 \frac{\sqrt{v_{\max} \log d}}{n}$ of Theorem 4.1. According to the well-known probabilistic guarantees for the Laplacian of the adjacency matrix of an ER model, $\lambda_2(\mathcal{I}^*) = \Theta_p\left(\frac{dp}{n}\right)$ and $v_{\max} = \Theta_p(dp)$ (see, e.g., Tropp, 2015). Thus, with high probability,

$$\begin{aligned} \lambda_2(\mathcal{I}^*) \geq c_0 \frac{\sqrt{v_{\max} \log d}}{n} &\Leftrightarrow \frac{dp}{n} = \Omega\left(\frac{\sqrt{dp \log d}}{n}\right) \\ &\Leftrightarrow p = \Omega\left(\frac{\log d}{d}\right), \end{aligned} \quad (13)$$

for some universal constant $c_0 > 0$. We mention in passing that, in this case, the requirements on $\lambda_2(\mathcal{I}^*)$ posed in Theorem 4.1 and Theorem 3.2 are equivalent.

As for the ℓ_2 error rates, our bound and the one in Chen et al. (2019) are as follows:

$$\begin{aligned} \text{Theorem 4.1:} \quad & \|h_d(\Delta)\|_2^2 \leq c_1 \frac{\kappa^2 t}{pL} \\ \text{Chen et al. (2019):} \quad & \|\Delta\|_2^2 \leq c' \frac{1}{pL} \end{aligned} \quad (14)$$

with probability at least $1 - e^{-t} - \frac{1}{d}$, for any $t > 0$, and $1 - O(d^{-7})$, respectively. Above, c' and κ are quantities depending on the dynamic range B . In the Erdős-Rényi graph case, $\kappa \leq (e^{-B} + e^B)^2$, where the dependence of c' on the dynamic range is implicit. The results in Eq. (14) are similar except for the use of the proxy function in Theorem 4.1. When the right hand side of the first inequality in Eq. (14) is bounded, $h_d(\Delta)$ can be lower bounded by $\Omega(\|\Delta\|_2^2)$, so that the both theorems yield essentially the

same ℓ_2 convergence rate. The fact that our general bound, applicable to arbitrary comparison graphs, yields the same convergence rate as Theorem 3.1 in Chen et al. (2019) is remarkable. Indeed, the analysis of Chen et al. (2019) is specifically tailored to the Erdős-Rényi graph, as the authors relied crucially on the special properties and the high degree of regularity of this type of graph.

Arbitrary graph under bounded dynamic range assumption. We now show that the error bound of Theorem 4.1 is similar to the one derived in Shah et al. (2016) under the same settings considered by the authors, namely an arbitrary comparison graph and bounded dynamic range B . It is worth repeating that our results concern the MLE of the BTL parameters and does not require knowledge of B , while Shah et al. (2016) make the additional convenient – and arguably impractical – assumption of a known, bounded dynamic range, which allows them to focus on the simpler constrained MLE.

In detail, Theorem 1 of Shah et al. (2016) analyze the ℓ_∞ regularized MLE

$$\tilde{w} := \arg \min_{w: \|w\|_\infty \leq B} l(w) \quad (15)$$

and provides the ℓ_2 -norm bound of the error $\tilde{\Delta} := \tilde{w} - w^*$ as

$$\|\tilde{\Delta}\|_L^2 \leq c'' \frac{\zeta^2 td}{\gamma^2 n}, \quad (16)$$

with probability at least $1 - e^{-t}$ for any $t > 0$ where c'' is an universal constant. Here, γ and ζ are parameters accounting for the curvature of the log-likelihood function and performance distribution across the items, defined by

$$\begin{aligned} \gamma &:= \min_{x \in [-2B, 2B]} \frac{d^2}{dt^2} (-\log \psi(x)) \\ \zeta &:= \max_{x \in [0, 2B]} \frac{\psi'(x)}{\psi(2x)(1-\psi(2x))}, \end{aligned} \quad (17)$$

and relate to the parameter κ in Eq. (12). Under the Bradley-Terry-Luce model where $\psi(x) = \frac{1}{1+e^{-x}}$, they are both functions of the dynamic range parameter B : $\gamma = \frac{1}{(e^{-B} + e^B)^2}$ and $\zeta = 1$. Then, Eq. (16) becomes

$$\|\tilde{\Delta}\|_L^2 \leq c'' (e^{-B} + e^B)^4 \frac{td}{n}. \quad (18)$$

On the other hand, repeating the same argument leading up to Eq. (8), we obtain that $\mathcal{I}^* \succeq \frac{M}{(1+M)^2} L$ where $M = e^{2B}$. Thus, under the bounded dynamic range, $\|h_d(\Delta)\|_{\mathcal{I}^*}^2 \geq \frac{1}{(e^{-B} + e^B)^2} \|h_d(\Delta)\|_L^2$ and also $\kappa \leq (e^{-B} + e^B)^2$. Thus, Theorem 4.1 yields that

$$\|h_d(\Delta)\|_L^2 \leq c_1 (e^{-B} + e^B)^4 \frac{td}{n} \quad (19)$$

with probability at least $1 - e^{-t} - \frac{1}{d}$ for any $t > 0$. The bounds (18) and (19) for the MLE and regularized MLE respectively (the latter requiring knowledge of the parameter

B) are identical except for a potential difference in the universal constants and the use of the proxy function in Eq. (19). As explained in the previous paragraph, the proxy function yields the same convergence rate as in Eq. (18) if the left hand side of Eq. (19) is bounded; otherwise, Theorem 4.1 provides a looser bound than Shah et al. (2016).

We also note that Hendrickx et al. (2019) established a BTL parameter estimator with minimax property in the sine error measure. Under bounded dynamic range, the sine error measure is equivalent to ℓ_2 error. Hence, the ℓ_2 estimation rate implied by Hendrickx et al. (2019) is strictly sharper than one implied by ours and Shah et al. (2016), which is expected. For example, in Section 4.1 therein, Shah et al. (2016) discussed the suboptimality of the ℓ_2 upperbound in their Theorem 2 under particular types of comparison graphs because the upperbound does not match its companion lowerbound. Instead, the sharper rate of the lowerbound is achieved by the ℓ_2 rate implication of Theorem 1, Hendrickx et al. (2019); compare, e.g., the examples in Section 4.1, Shah et al. (2016) and Table 1, Hendrickx et al. (2019). For the proper comparison, the sine error rate in the table should be squared and multiplied by “ $\frac{n^2}{k}$ ” in their notation. (Note that “ n ” is the number of the compared items and “ k ” is the number of comparison in Hendrickx et al. (2019).) On the other hand, the estimation rates of Shah et al. (2016) are minimax in the $\|\cdot\|_L$ error. Because our bound inherits the minimax property of Shah et al. (2016) under bounded dynamic range, it is generally better than Hendrickx et al. (2019) in the $\|\cdot\|_L$ error.

In the present discussion, since we are assuming a general comparison graph, we have bounded both the ratio ζ/γ and κ , appearing in (18) and (19) respectively, using B . While ζ/γ are always functions of B only regardless of the comparison graph, the parameter κ also depends on the specific comparison graph topology. This dependence can in turn be leveraged in some specific cases to reduce the value of κ and thus produce shaper rates via Theorem 4.1. We elaborate on this point next.

Banded graph without bounded dynamic range assumption. Theorem 4.2 of Shah et al. (2016) and Theorem 4.1 differ crucially in the ways they account in the error rate for the dependence on the gap in performance among items and on the comparison graph topology. Theorem 4.2 of Shah et al. (2016) decouples these two aspects. The gap in performance enters the ℓ_2 bound (18) only through the ratio ζ/γ , which depends on the dynamic range parameter B but not the comparison graph; on the other hand, the graph topology affects the rate only through the algebraic connectivity $\lambda_2(L)$, which is independent of the gap in performance. In contrast, the ℓ_2 rate in Theorem 4.1 depends on κ and $\lambda_2(\mathcal{I}^*)$, both of which simultaneously quantify the impact of the graph topology and the gap in performance,

and their potential interaction. When we crudely decouple such dependence, as we did above, such interaction is lost and we essentially recover the bound in Theorem 4.2 of Shah et al. (2016). But a more careful use of the dependence will lead to better results by Theorem 4.1. We illustrate this fact using a banded graph topology.

Suppose that w_i^* 's are evenly distributed along the dynamic range

$$w_i^* := \frac{2i-d}{d}B, \quad i \in [d] \quad (20)$$

and that the comparisons are made T times only between items having a difference in ranking smaller than a fixed positive numbers $W > 0$. The resulting normalized graph Laplacian is

$$L_{ij} = \begin{cases} -\sum_{k:k \neq i} L_{ik}, & i = j, \\ -1/n, & 0 < |i-j| \leq W, \\ 0, & \text{elsewhere,} \end{cases} \quad (21)$$

and the total number of comparisons n is $T(dW - W(W-1)/2)$. Since the graph Laplacian is a banded matrix, we call the graph a *banded comparison graph* and the hyperparameter W the *comparison width*. Banded comparison graphs are easily seen in real data, in particular, when the size of the item pool is huge. In the high-dimensional setting, it is impossible to compare every pair of the items due to financial and physical constraints. To maximize the efficiency of limited resource, comparisons are often made only between the items of which none are overwhelming the other. The influence of this comparison graph to the estimation error is measured by its algebraic connectivity, which we establish in the following lemma. See Appendix A.3.1 for the proof.

Lemma 4.2. *The algebraic connectivity of the given banded comparison graph is $\lambda_2(L) = \Theta\left(\frac{W^2}{d^3}\right)$.*

Since γ and ζ for Theorem 2 in Shah et al. (2016) depend only on the dynamic range, their values are still as in (17), and the resulting ℓ_2 error bound from Theorem 4.2 of Shah et al. (2016) is

$$\|\tilde{\Delta}\|_2^2 \leq c'' \frac{\zeta^2}{\gamma^2 \lambda_2(L)} \frac{td}{n} \leq 2c'' \frac{td^3}{TW^3} e^{4B}, \quad (22)$$

with probability at least $1 - e^{-t}$ for any $t > 0$ and for some the universal constant c'' (See Eq. (18)).

On the other hand, Theorem 4.1 leads to a tighter bound. To see this, the Fisher information matrix for this graph topology is

$$\mathcal{I}_{ij}^* = \begin{cases} -\sum_{k:k \neq i} \mathcal{I}_{ij}^*, & i = j, \\ -\frac{1/(d-1)}{(e^{-\frac{B}{d}|i-j|} + e^{\frac{B}{d}|i-j|})^2}, & 0 < |i-j| \leq W, \\ 0, & \text{elsewhere,} \end{cases} \quad (23)$$

where $|\mathcal{I}_{ij}(w^*)| \geq \frac{1}{4e^{2BW/d}} |L_{ij}|$ for every (i, j) such that $0 < |i - j| \leq W$. As a result, κ and $\lambda_2(\mathcal{I}^*)$ turns out to be much smaller than in the previous example: $\kappa \leq 4e^{2BW/d}$ and $\lambda_2(\mathcal{I}^*) = \Omega\left(\frac{1}{d^3 e^{2BW/d}}\right)$. Putting all these pieces together, we arrive at a sharper ℓ_2 error bound

$$\|h_d(\Delta)\|_2^2 \leq c_1 \frac{\kappa}{\lambda_2(\mathcal{I}^*)} \frac{td}{n} \leq c_1 \frac{td^3}{TW^3} e^{\frac{4BW}{d}} \quad (24)$$

with probability at least $1 - \frac{1}{d} - e^{-t}$.

When B increases in d , Eq. (24) provides a much tighter convergence rate so the ℓ_2 convergence requires a much smaller sample complexity in terms of d . Concretely, suppose that $\|w^*\|_\infty$ increases at a rate $\sqrt{\log d}$ – this would be for example the case if the w_i^* 's were independently sampled from a standard Gaussian distribution, with high probability. Then, Theorem 2 of (Shah et al., 2016) implies that every compared pair should be matched $T = \omega(e^{8\sqrt{\log d}} \log d)$ times for ℓ_2 consistency even if $W = d$. On the other hand, Theorem 4.1 suggests a much milder condition that every compared pair should be matched $T = \omega((\log d)^{5/2})$ times when $W = \Theta\left(\frac{d}{\sqrt{\log d}}\right)$.

We remark that this result does not contradict the minimax lower bound found in Theorem 1 of Shah et al. (2016). In fact, although their upper bound is proved to be optimal with respect to d and n under fixed B , it has a sub-optimal dependence on B at least when compared with the minimax lower bound therein. In the above example, Theorem 4.1 yields improvements in the dependence to B and graph topology by incorporating them into Fisher information matrix \mathcal{I}^* . In Section 5, we demonstrate the improved dependence in Theorem 4.1 by simulations, particularly with respect to the dynamic range and graph topology.

5. Illustrative Simulations

Shah et al. (2016) demonstrated the tightness of the error bounds in their Theorem 2 therein with respect to the number of the comparisons n through simulations. We do the same for Theorem 4.1 but with respect to the dynamic range and the graph topology. The comparison data were simulated using the BTL model (Eq. (2)) under the banded comparison graphs introduced in Section 4.1. In our experiment, the number of compared items d and the number of comparisons between compared pairs T are fixed to 100 and 5, respectively. The maximal performance gap B ranges from $4e^{-1}\sqrt{\log d}$ to $4e\sqrt{\log d}$, and the comparison width W from $e^{-1}\frac{d}{4\sqrt{\log d}}$ to $e\frac{d}{4\sqrt{\log d}}$. The MLE was calculated by iterative Luce Spectral Ranking (Maystre & Grossglauser, 2015), provided by Python package `choix`¹. For each pair of values (B, W) , we performed 100 simulations, and the average of the estimated ℓ_2 error is plotted




Figure 2. Estimated ℓ_2 error by simulation compared to the theorems. ℓ_2 rates are estimated under various (a) maximal performance differences B and (b) graph topologies induced by comparison widths W . With log-scaled y-axes, the blue lines indicate the averages of the ℓ_2 errors over the 100 simulation runs where the blue shaded areas represent the 95% pointwise confidence interval for the averages. The yellow and green lines indicate the upper bounds of the ℓ_2 errors suggested by Theorem 4.1 and Theorem 2 of Shah et al. (2016), respectively, up to constant scales (which are constant shifts in the plots). The zoom pane only shows the estimated values and the suggested upperbound by Theorem 4.1 for the readability. With respect to both B and W , Theorem 4.1 gives tighter bounds to the ℓ_2 error.

in Fig. 2 (blue line) along with the consistency rates implied by Theorem 4.1 (yellow line) and Theorem 2 of Shah et al. (2016) (green line). The light blue shaded area represents 95% pointwise confidence intervals for the estimated averages. Fig. 2a shows the simulation results when W was fixed at $\frac{d}{\sqrt{\log d}}$ and B changed across the range while Fig. 2b displays the result for B at the fixed value of $\sqrt{\log d}$ and W changing. Because Theorem 4.1 and Theorem 2 of Shah et al. (2016) both provide ℓ_2 consistency rates up to a universal scale, the yellow and green lines in Fig. 2 are arbitrarily shifted to make the comparison easier. Hence, when comparing our bound (orange line) with that of (Shah et al., 2016) (green line), we should not compare the relative position of the corresponding curves but their rates of increase as B increases or W decreases, especially in relation to the blue curve, based on the data. The plots clearly indicate that Theorem 4.1 provides tighter rates in both B and W . We obtained similar results for varying B in other simulations involving different graph topologies. See Appendix B and github.com/HeejongBong/mmpc for the supplementary simulation results and reproducible

¹Published through PyPi under MIT License

simulation code scripts.

6. Discussion

We have derived novel conditions for the existence of the MLE and novel ℓ_2 estimation rates in the BTL model under weaker assumptions than those assumed in the literature so far. In particular, we allow arbitrary comparison graphs and do not require a bounded gap among the quality scores, nor knowledge of such bound. Our bounds, based on the Fisher information matrix, recover existing results as corollaries and are also applicable to more general settings of practical relevance. Our theoretical results support the use of the BTL model for ranking problems and more generally statistical inference with pairwise comparison data.

To the best of our knowledge, the use of proxy functions in analyzing the MLE solutions and their properties is new, although convex surrogate loss functions are widely used in classification. Because the BTL model is a special case of the logistic regression, one might use the same technique to establish a tight uncertainty measure for general logistic models when the odds are not necessarily bounded. For example, a promising open problem is the sine error minimax estimate of the BTL model without bounded dynamic range by incorporating the proxy function technique and Hendrickx et al. (2019).

Another important topic in the ranking literature which we did not discuss here is the probability bound of identifying the K best items from the observed partial orderings: namely, top- K ranking. Recently, the minimax bound under the Erdős-Rényi comparison graph was established by Chen et al. (2020a; 2019), and a decision-theoretic recovery of full ranking was suggested by Chen et al. (2020b). However, the learning theory of top- K ranking from passive observation of pair-wise comparisons in arbitrary graphs is yet unknown. As discussed in Chen et al. (2020a), theoretical guarantee for top- K recovery requires the ℓ_∞ estimation error to be controlled. To establish such bound, the performance score w^* and comparison graph L should be more fine-tuned.

One might point out that the banded comparison graphs, used for the comparison against Shah et al. (2016), already contained prior knowledge about w^* , which is not properly evaluated in Section 4.1. However, such comparison graphs can emerge without prior knowledge in online settings where the comparison graph is decided based on the recent estimate of the BTL parameters. A promising direction of future study is to study the theoretical property of such online estimators.

Acknowledgements

We gratefully acknowledge helpful comments and mindful suggestions from the reviewers and meta-reviewer. We also send the deepest appreciation for inspiring discussions with Wanshan Li and Shamindra Shrotriya in the Department of Statistics and Data Sciences, Carnegie Mellon University.

References

- 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.
- Brouwer, A. E. and Haemers, W. H. *Spectra of graphs*. Springer Science & Business Media, 2011.
- Chen, P., Gao, C., and Zhang, A. Y. Partial recovery for top- k ranking: Optimality of mle and sub-optimality of spectral method. *arXiv preprint arXiv:2006.16485*, 2020a.
- Chen, P., Gao, C., and Zhang, A. Y. Optimal full ranking from pairwise comparisons. *arXiv preprint arXiv:2101.08421*, 2020b.
- Chen, Y., Fan, J., Ma, C., and Wang, K. Spectral method and regularized MLE are both optimal for top- K ranking. *Ann. Statist.*, 47(4):2204–2235, 2019. ISSN 0090-5364. doi: 10.1214/18-AOS1745.
- Erdős, P. and Rényi, A. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5(1):17–60, 1960.
- Ford Jr, L. R. Solution of a ranking problem from binary comparisons. *The American Mathematical Monthly*, 64(8P2):28–33, 1957.
- Hajek, B., Oh, S., and Xu, J. Minimax-optimal inference from partial rankings. *Advances in Neural Information Processing Systems*, 27, 2014.
- Han, R., Ye, R., Tan, C., Chen, K., et al. Asymptotic theory of sparse bradley–terry model. *Annals of Applied Probability*, 30(5):2491–2515, 2020.
- Hanson, D. L. and Wright, F. T. A bound on tail probabilities for quadratic forms in independent random variables. *The Annals of Mathematical Statistics*, 42(3):1079–1083, 1971.
- Hendrickx, J., Olshevsky, A., and Saligrama, V. Graph resistance and learning from pairwise comparisons. In *International Conference on Machine Learning*, pp. 2702–2711. PMLR, 2019.
- Hendrickx, J., Olshevsky, A., and Saligrama, V. Minimax rate for learning from pairwise comparisons in the btl model. In *International Conference on Machine Learning*, pp. 4193–4202. PMLR, 2020.

- Khetan, A. and Oh, S. Computational and statistical trade-offs in learning to rank. *Advances in Neural Information Processing Systems*, 29, 2016.
- Kolokolnikov, T., Ousting, B., and Von Brecht, J. Algebraic connectivity of erdős-rényi graphs near the connectivity threshold. 2014.
- Luce, R. D. *Individual choice behavior: A theoretical analysis*. Courier Corporation, 2012.
- Maystre, L. and Grossglauser, M. Fast and accurate inference of plackett-luce models. Technical report, 2015.
- Negahban, S., Oh, S., and Shah, D. Iterative ranking from pair-wise comparisons. *Advances in neural information processing systems*, 25:2474–2482, 2012.
- 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.
- Simons, G. and Yao, Y.-C. Asymptotics when the number of parameters tends to infinity in the bradley-terry model for paired comparisons. *The Annals of Statistics*, 27(3): 1041–1060, 1999.
- Tropp, J. A. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8: 1–230, 2015.
- Wainwright, M. J. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press, 2019.
- Yan, T., Yang, Y., and Xu, J. Sparse paired comparisons in the bradley-terry model. *Statistica Sinica*, pp. 1305–1318, 2012.
- Zermelo, E. Die Berechnung der Turnier-Ergebnisse als ein Maximumproblem der Wahrscheinlichkeitsrechnung. *Math. Z.*, 29(1):436–460, 1929.

A. Proofs

In the following proofs, n_{ij} denotes the number of comparisons between i and j . Accordingly, the Laplacian of the comparison graph consists of elements

$$L_{ij} = \begin{cases} -\frac{n_{ij}}{n}, & i \neq j, \\ -\sum_{k:k \neq i} L_{ik}, & i = j. \end{cases} \quad (25)$$

Denoting by p_{ij} the winning probability of i against j for $i, j = 1, \dots, d$, we model the number of comparisons where i defeated j by a Binomial random variable $A_{ij} \sim \text{Binom}(n_{ij}, p_{ij})$. Then, by straightforward calculations,

$$l(w) := -\frac{1}{n} \sum_{i,j:i>j} \{A_{ij} \log(1 + e^{w_j - w_i}) + A_{ji} \log(1 + e^{w_i - w_j})\}, \quad (26)$$

$$[\nabla l(w^*)]_i = \frac{1}{n} \sum_{j:j \neq i} (A_{ij} - n_{ij} p_{ij}), \quad (27)$$

and the Fisher information matrix (the Hessian of the negative log-likelihood) at w^* has (i, j) coordinates of the form

$$[\mathcal{I}(w^*)]_{ij} = \begin{cases} -\frac{n_{ij} p_{ij} p_{ji}}{n}, & i \neq j, \\ -\sum_{k:k \neq i} [\mathcal{I}(w^*)]_{ik}, & i = j \end{cases} \quad (28)$$

for $i, j \in [d]$.

A.1. Proof for Theorem 3.2

For each non-empty $I \subset [d]$, we denote by E_I the event that none of the items in $I^c := [d] \setminus I$ has defeated any item in I . That is,

$$E_I = \{A_{ij} = n_{ij}, \forall i \in I, j \in I^c\}. \quad (29)$$

Since n_{ij} is the maximum value that each A_{ij} can take, the event E_I can be equivalently expressed as

$$E_I = \left\{ \sum_{i \in I, j \in I^c} A_{ij} = \sum_{i \in I, j \in I^c} n_{ij} \right\}. \quad (30)$$

Hence,

$$\begin{aligned} \mathbb{P}[E_I] &= \mathbb{P} \left[\sum_{i \in I, j \in I^c} A_{ij} = \sum_{i \in I, j \in I^c} n_{ij} \right] \\ &= \mathbb{P} \left[\sum_{i \in I, j \in I^c} (A_{ij} - n_{ij} p_{ij}) = \sum_{i \in I, j \in I^c} n_{ij} (1 - p_{ij}) \right] \\ &\leq \mathbb{P} \left[\sum_{i \in I, j \in I^c} (A_{ij} - n_{ij} p_{ij}) \geq \sum_{i \in I, j \in I^c} n_{ij} p_{ij} (1 - p_{ij}) \right]. \end{aligned} \quad (31)$$

Because the A_{ij} 's are independent Binomial random variables, we use Bernstein's inequality to control the probability of the event E_I :

$$\begin{aligned} \mathbb{P}[E_I] &\leq \mathbb{P} \left[\sum_{i \in I, j \in I^c} (A_{ij} - n_{ij} p_{ij}) \geq \sum_{i \in I, j \in I^c} n_{ij} p_{ij} (1 - p_{ij}) \right] \\ &\leq \exp \left(-\frac{3}{2} \sum_{i \in I, j \in I^c} n_{ij} p_{ij} (1 - p_{ij}) \right). \end{aligned} \quad (32)$$

The above bound can be expressed in terms of the Fisher information matrix $\mathcal{I}(w^*)$. To see this, by the formula Eq. (28) of the Fisher information matrix, we have that, for each $i \in I$ and $j \in I^c$, $n_{ij}p_{ij}(1-p_{ij}) = -n\mathcal{I}_{ij}(w^*)$. Next, letting e_I be the indicator vector of the subset I (i.e., the i -th element of e_I is 1 for $i \in I$ and 0 elsewhere) and noting that $e_{I^c} = \mathbf{1} - e_I$, we obtain that

$$\sum_{i \in I, j \in I^c} n_{ij}p_{ij}(1-p_{ij}) = -e_I \mathcal{I}(w^*) e_{I^c} = -e_I \mathcal{I}(w^*) (\mathbf{1} - e_I) = e_I \mathcal{I}(w^*) e_I,$$

where the last identity follows the fact that $\mathbf{1}$ is in the null space of $\mathcal{I}(w^*)$. This is because the comparison graph is connected and therefore $\mathcal{I}(w^*)$ has one zero eigenvalue with $\text{span}(\mathbf{1})$ as the associated eigenspace. Thus, we conclude that

$$\mathbb{P}[E_I] \leq \exp\left(-\frac{3n}{2} e_I^\top \mathcal{I}(w^*) e_I\right). \quad (33)$$

Next, using again the fact that the null space of $\mathcal{I}(w^*)$ is spanned by $\mathbf{1}$,

$$e_I^\top \mathcal{I}(w^*) e_I \geq \lambda_2(\mathcal{I}(w^*)) \|\text{proj}_{\text{span}(\mathbf{1})^\perp}(e_I)\|_2^2 = \lambda_2(\mathcal{I}(w^*)) \frac{|I|(d-|I|)}{d}, \quad (34)$$

where $\text{proj}_{\text{span}(\mathbf{1})^\perp}(\cdot)$ is the projection mapping onto the linear subspace of \mathbb{R}^d orthogonal to $\text{span}(\mathbf{1})$.

In the last step of the proof, we follow the arguments from Lemma 1 in (Simons & Yao, 1999). By the union bound,

$$\begin{aligned} \mathbb{P}[\text{Condition 3.1 fails}] &\leq \sum_{I \in 2^{[d]}} \exp\left(-\frac{3n}{2d} \lambda_2(\mathcal{I}(w^*)) |I|(d-|I|)\right) \\ &\leq 2 \sum_{i=1}^{\lceil d/2 \rceil} \binom{d}{i} \exp\left(-\frac{3n}{2d} \lambda_2(\mathcal{I}(w^*)) i(d-i)\right) \\ &\leq 2 \sum_{i=1}^{\lceil d/2 \rceil} \binom{d}{i} \exp\left(-\frac{3n}{2d} \lambda_2(\mathcal{I}(w^*)) i \frac{d}{2}\right) \\ &\leq 2 \left[\left(1 + \exp\left(-\frac{3n}{4} \lambda_2(\mathcal{I}(w^*))\right)\right)^d - 1 \right]. \end{aligned} \quad (35)$$

The desired result in Equation (7) is a straightforward result after plugging-in $\lambda_2(\mathcal{I}^*) \geq 2 \frac{\log d}{n}$.

A.2. Proof for Theorem 4.1

Let $l^*(w)$ be the expected log-likelihood function at $w \in \mathbb{R}^d$ under the groundtruth score parameter w^* , i.e.

$$l^*(w) := -\frac{1}{n} \sum_{i,j:i>j} \{n_{ij}p_{ij} \log(1 + e^{w_j - w_i}) + n_{ij}p_{ji} \log(1 + e^{w_i - w_j})\} \quad (36)$$

Since $\log(1 + e^{w_j - w_i}) = (w_j - w_i) + \log(1 + e^{w_i - w_j})$, we have that

$$l^*(w) := -\frac{1}{n} \sum_{i,j:i>j} \{n_{ij}p_{ij}(w_j - w_i) + n_{ij} \log(1 + e^{w_i - w_j})\}, \quad (37)$$

and

$$l(w) := -\frac{1}{n} \sum_{i,j:i>j} \{A_{ij}(w_j - w_i) + n_{ij} \log(1 + e^{w_i - w_j})\}. \quad (38)$$

Thus,

$$\Delta l(w) := l(w) - l^*(w) = -\frac{1}{n} \sum_{i,j:i>j} (A_{ij} - n_{ij}p_{ij})(w_j - w_i) = w^\top \nabla l(w^*). \quad (39)$$

Now, let W denote the set of parameter values for which the log-likelihood function is higher than at the true score parameter, i.e., $W := \{w \in \mathbb{R}^d : \mathbf{1}^\top w = 0, l(w) \geq l(w^*)\}$. Since the MLE is one of the elements of W , if W is ℓ_2 -bounded by some

$r > 0$, then the MLE also exists and is ℓ_2 -bounded by $r > 0$. Throughout the rest of this proof, we let \hat{w} be an arbitrary element in W (possibly the MLE), and $\Delta := \hat{w} - w^*$. Due to the higher log-likelihood condition of W ,

$$\begin{aligned} l^*(w^*) - l^*(\hat{w}) &\leq (l^*(w^*) - l^*(\hat{w})) + (l(\hat{w}) - l(w^*)) \\ &\leq \Delta l(\hat{w}) - \Delta l(w^*) = \Delta^\top \nabla l(w^*). \end{aligned} \quad (40)$$

On the other hand,

$$\begin{aligned} &l^*(w^*) - l^*(\hat{w}) \\ &= \frac{1}{n} \sum_{i,j:i>j} \left\{ n_{ij} p_{ij} \log \left(\frac{1 + e^{\hat{w}_j - \hat{w}_i}}{1 + e^{w_j^* - w_i^*}} \right) + n_{ij} p_{ji} \log \left(\frac{1 + e^{\hat{w}_i - \hat{w}_j}}{1 + e^{w_i^* - w_j^*}} \right) \right\} \\ &= \frac{1}{n} \sum_{i,j:i>j} n_{ij} p_{ij} p_{ji} f(w_i^* - w_j^*, \Delta_i - \Delta_j) \\ &= \sum_{i,j:i>j} -[\mathcal{I}(w^*)]_{ij} f(w_i^* - w_j^*, \Delta_i - \Delta_j) \end{aligned} \quad (41)$$

where

$$f(x, y) = (1 + e^{-x}) \log \left(\frac{1 + e^{x+y}}{1 + e^x} \right) + (1 + e^x) \log \left(\frac{1 + e^{-x-y}}{1 + e^{-x}} \right). \quad (42)$$

We observe that the curvature of $f(w_i^* - w_j^*, \Delta_i - \Delta_j)$ with respect to Δ converges to 0 as $\Delta \rightarrow \infty$, and this property prevents $l^*(w^*) - l^*(\hat{w})$ from being strongly convex. We bypass this problem by deploying the proxy function h_d as defined in Section 4 and show the strong convexity of $l^*(w^*) - l^*(\hat{w})$ and $f(w_i^* - w_j^*, \Delta_i - \Delta_j)$ with respect to $h_d(\Delta)$. First, we show in the next result that $f(w_i^* - w_j^*, \Delta_i - \Delta_j)$ is lower bounded by a quadratic function of $h_d(\Delta_i)$; see Appendix A.3.2 for the proof.

Lemma A.1. $f(w_i^* - w_j^*, \Delta_i - \Delta_j) \geq c_a (h(\Delta_i) - h(\Delta_j))^2$ for some universal constant $c_a > 0$.

Thus,

$$\begin{aligned} l^*(w^*) - l^*(\hat{w}) &= \sum_{i,j:i>j} -[\mathcal{I}(w^*)]_{ij} f(w_i^* - w_j^*, \Delta_i - \Delta_j) \\ &\geq c_a h_d(\Delta)^\top \mathcal{I}(w^*) h_d(\Delta), \end{aligned} \quad (43)$$

which implies strong convexity of $l^*(w^*) - l^*(\hat{w})$ with respect to $h_d(\Delta)$. Using this, we arrive at the key basic inequality

$$\Delta^\top \nabla l(w^*) \geq c_a h_d(\Delta)^\top \mathcal{I}(w^*) h_d(\Delta). \quad (44)$$

We prove Theorem 4.1 starting from the basic inequality and further relying on the following, easily verifiable but special property of h :

$$\text{sgn}(\Delta_i) h(\Delta_i)^2 = \Delta_i - 2h(\Delta_i), \quad \text{for } i = 1, \dots, d. \quad (45)$$

From this, we arrive at the following, nontrivial result; see Appendix A.3 for the proof.

Lemma A.2. *Provided that $\sum_i \Delta_i = 0$,*

$$\sum_i (h(\Delta_i))^2 \geq \frac{2}{d} \left(\sum_i h(\Delta_i) \right)^2. \quad (46)$$

With the help of the lemma, we obtain that

$$\|h_d(\Delta)\|_2^2 \leq 2\|h(\Delta)\|_2^2 - \frac{2}{d} \left(\sum_i h(\Delta_i) \right)^2 = 2h_d(\Delta)^\top \left(I - \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right) h_d(\Delta). \quad (47)$$

Plugging Eq. (45) into the basic inequality Eq. (44), we get that

$$h_d(\Delta)^\top \mathcal{I}(w^*) h_d(\Delta) \leq c_a \left\{ 2h_d(\Delta)^\top \nabla l(w^*) + \sum_{i=1}^d \text{sgn}(\Delta_i) h(\Delta_i)^2 \nabla_i l(w^*) \right\}. \quad (48)$$

Hence,

$$\begin{aligned} h_d(\Delta)^\top \mathcal{I}(w^*) h_d(\Delta) &\leq c_a \left\{ 2h_d(\Delta)^\top \nabla l(w^*) + \|\nabla l(w^*)\|_\infty \|h_d(\Delta)\|_2^2 \right\} \\ &\leq 2c_a \left\{ h_d(\Delta)^\top \nabla l_d(w^*) + \|\nabla l(w^*)\|_\infty h_d(\Delta)^\top \left(I - \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right) h_d(\Delta) \right\}, \end{aligned} \quad (49)$$

where the last step is an application of Eq. (47). Moving the quadratic terms to the left-hand side,

$$h_d(\Delta)^\top \left(\mathcal{I}(w^*) - 2c_a \|\nabla l(w^*)\|_\infty \left(I - \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right) \right) h_d(\Delta) \leq 2c_a h_d(\Delta)^\top \nabla l(w^*). \quad (50)$$

Now, suppose that it holds that

$$\mathcal{I}(w^*) \succeq 4c_a \|\nabla l(w^*)\|_\infty \left(I - \frac{1}{d} \mathbf{1}\mathbf{1}^\top \right) \quad (51)$$

with high probability; we will prove this fact below. Then, taking this as a given, we have that

$$\frac{1}{2} h_d(\Delta)^\top \mathcal{I}(w^*) h_d(\Delta) \leq 2c_a h_d(\Delta)^\top \nabla l, \quad (52)$$

and, by Lemma 9 of Shah et al. (2016),

$$\frac{1}{2} \|h_d(\Delta)\|_{\mathcal{I}(w^*)}^2 \leq 2c_a \|h_d(\Delta)\|_{\mathcal{I}(w^*)} \|\nabla l(w^*)\|_{\mathcal{I}(w^*)+}. \quad (53)$$

As a result,

$$\|h_d(\Delta)\|_{\mathcal{I}(w^*)} \leq 4c_a \|\nabla l(w^*)\|_{\mathcal{I}(w^*)+} \leq 4c_a \sqrt{\kappa} \|\nabla l(w^*)\|_{L^+}, \quad (54)$$

where the last step is straightforward from the definition of κ in Theorem 4.1. Using the Hanson-Wright inequality (Hanson & Wright, 1971) as in the proof of Theorem 1 of (Shah et al., 2016), we obtain that

$$\|\nabla l(w^*)\|_{L^+}^2 \leq \frac{td}{n},$$

with probability at least $1 - e^{-t}$. Putting everything together, we thus finally have that

$$\|h_d(\Delta)\|_{\mathcal{I}(w^*)}^2 \leq c_1 \kappa \frac{td}{n} \quad (55)$$

with probability at least $1 - e^{-t} - p$ where p is the probability that Eq. (51) does not hold.

To complete the proof, it remains to prove that Eq. (51) holds with high probability. To that effect, we have that, for each $i \in [d]$,

$$[\nabla l(w^*)]_i = \frac{1}{n} \sum_{j:j \neq i} \left(x_{ij} - \frac{x_{ij} + x_{ji}}{1 + e^{w_j^* - w_i^*}} \right) \quad (56)$$

is a sub-Gaussian random variable with $\sigma^2 = v_{\max}/4$. By a standard maximal inequality for sub-Gaussian random variables (Wainwright, 2019),

$$\mathbb{P} \left[\|\nabla l(w^*)\|_\infty \geq \frac{\sqrt{v_{\max}}}{2n} (\sqrt{2 \log(2d)} + t) \right] \leq e^{-t^2/2} \quad (57)$$

for any $t > 0$. Setting $t = \sqrt{2 \log(2d)}$, we have that

$$\mathbb{P} \left[\|\nabla l(w^*)\|_\infty \geq \frac{1}{n} \sqrt{2v_{\max} \log(2d)} \right] \leq \frac{1}{2d}. \quad (58)$$

Hence, provided that $\lambda_2(\mathcal{I}(w^*)) \geq \frac{c_0}{n} \sqrt{v_{\max} \log d}$ for some large enough $c_0 > 0$, $\lambda_2(\mathcal{I}(w^*)) \geq \|\nabla l(w^*)\|_\infty$ with probability at least $1 - \frac{1}{2d}$, for sufficiently large d . This follows by a straightforward derivation of Eq. (51).

A.3. Proofs of Lemmas

In this section we provide the proofs of technical Lemmas needed in the proof of Theorem 4.1.

A.3.1. PROOF OF LEMMA 4.2

In this proof, we use that this graph has a similar algebraic connectivity to the Cayley graph on $(\mathbb{Z}_d, +)$ with difference set $\{-W, \dots, -1\} \cup \{1, \dots, W\}$. This is formalized in the following argument, which leverages results from the spectral graph theory for the Cayley graph (e.g., [Brouwer & Haemers, 2011](#))

Let \tilde{L} be the normalized graph Laplacian of the Cayley graph on $(\mathbb{Z}_d, +)$ with difference set $\{-W, \dots, -1\} \cup \{1, \dots, W\}$. \tilde{L} has elements

$$\tilde{L}_{ij} = \begin{cases} -\sum_{k:k \neq i} \tilde{L}_{ik}, & i = j, \\ -1/\tilde{n}, & 0 < |i - j| \leq W, \\ -1/\tilde{n}, & 0 < d - |i - j| \leq W, \\ 0, & \text{elsewhere,} \end{cases} \quad (59)$$

where $\tilde{n} = dW$. It follows from Proposition 1.7.1 in [Brouwer & Haemers \(2011\)](#) that $\lambda_2(\tilde{L}) \geq \lambda_2(L)$.

On the other hand, let \bar{L} be the normalized graph Laplacian of the Cayley graph on $(\mathbb{Z}_{2d}, +)$ with the same difference set. \bar{L} can be obtained from Eq. (59) by changing d to $2d$. Suppose that we fold this graph so that vertices i and $2d - i$ are equated. The obtained folded graph has normalized graph Laplacian L' such that

$$L'_{ij} = \begin{cases} -\sum_{k:k \neq i} L'_{ik}, & i = j, \\ -3/n', & i \neq j \text{ and } i + j \leq W + 2, \\ -3/n', & i \neq j \text{ and } i + j \geq d - W, \\ 0, & |i - j| > W, \\ -1/n', & \text{elsewhere,} \end{cases} \quad (60)$$

where $n' = 2dW$. If (x_1, \dots, x_d) is an eigenvector of L' with eigenvalue λ , then $(x_1, \dots, x_d, x_d, \dots, x_1)$ is an eigenvector of \bar{L} with the same eigenvalue. Hence, $\lambda_2(\bar{L})$ is always smaller than $\lambda_2(L')$. Now, we notice that $L'_{ij} \leq 3L_{ij}$ for every $(i, j) \in [d]^2$ and hence Proposition 1.7.1 in [Brouwer & Haemers \(2011\)](#) again implies $\lambda_2(L') \leq 3\lambda_2(L)$. Putting everything together,

$$\frac{1}{3}\lambda_2(\bar{L}) \leq \frac{1}{3}\lambda_2(L') \leq \lambda_2(L) \leq \lambda_2(\tilde{L}). \quad (61)$$

According to the spectral properties of Cayley graphs listed in [Brouwer & Haemers \(2011\)](#),

$$\begin{aligned} \lambda_2(\tilde{L}) &\approx \frac{2W}{n} - \frac{2}{n} \sum_{i=1}^W \cos\left(\frac{2\pi i}{d}\right) \\ &= \frac{2}{n} \sum_{i=1}^W \left(1 - \cos\left(\frac{2\pi i}{d}\right)\right) \\ &= \Theta\left(\frac{d}{n\pi} \int_0^{2\pi W/d} (1 - \cos x) dx\right) \\ &= \Theta\left(\frac{W^3}{nd^2}\right). \end{aligned} \quad (62)$$

Similarly, $\lambda_2(\bar{L}) = \Theta\left(\frac{W^2}{d^3}\right)$, and $\lambda_2(L)$ has the same rate in d .

A.3.2. PROOF OF LEMMA A.1

Let

$$g(y) := \log\left(\frac{1 + e^y}{2}\right) + \log\left(\frac{1 + e^{-y}}{2}\right). \quad (63)$$

We will derive the desired inequality from the following two properties of f , g , and h :

- $f(x, y) \geq g(y)$ for any $x, y \in \mathbb{R}$, and
- $g(y_1 - y_2) \geq c_a (h(y_1) - h(y_2))^2, \forall y_1, y_2 \in \mathbb{R}$ for some universal constant c_a .

First, for any $x, y \in \mathbb{R}$,

$$\frac{\partial}{\partial y} f(x, y) = (1 + e^{-x}) \frac{e^{x+y}}{1 + e^{x+y}} - (1 + e^x) \frac{e^{-x-y}}{1 + e^{-x-y}} \quad (64)$$

and

$$\frac{\partial}{\partial y} g(y) = \frac{e^y}{1 + e^y} - \frac{e^{-y}}{1 + e^{-y}}. \quad (65)$$

Hence,

$$\begin{aligned} \frac{\partial}{\partial y} (f(x, y) - g(y)) &= \frac{e^x + e^y}{(1 + e^{x+y})(1 + e^y)(1 + e^{-y})} (e^y - e^{-y}) \\ &\begin{cases} > 0, & y > 0, \\ = 0, & y = 0, \\ < 0, & y < 0. \end{cases} \end{aligned} \quad (66)$$

With the fact that $\frac{\partial}{\partial y} (f(x, y) - g(y)) \big|_{y=0} = 0$, we derive the first inequality between f and g .

Now, we move on to the second inequality between g and h . We first note that

$$\sqrt{x+h} - \sqrt{x} \geq \sqrt{y+h} - \sqrt{y} \quad (67)$$

where $0 < x < y$ and $h > 0$, which can be easily seen through the following inequality:

$$\sqrt{x+h} - \sqrt{x} = \frac{h}{\sqrt{x+h} + \sqrt{x}} \geq \frac{h}{\sqrt{y+h} + \sqrt{y}} = \sqrt{y+h} - \sqrt{y}. \quad (68)$$

Hence,

$$(h(y_1) - h(y_2))^2 = (\sqrt{|y_1|+1} - \sqrt{|y_2|+1})^2 \leq (\sqrt{||y_1| - |y_2|| + 1} - 1)^2 \leq h(y_1 - y_2)^2. \quad (69)$$

That is, it suffices to show that $g(x) \geq c_a h(x)^2, \forall x$ for some universal constant c_a ; then,

$$c_a (h(y_1) - h(y_2))^2 \leq c_a h(y_1 - y_2)^2 \leq g(y_1 - y_2). \quad (70)$$

We note that both $g(\cdot)$ and $h(\cdot)^2$ increase linearly around 0 and so do $\sqrt{g(\cdot)}$ and $h(\cdot)$ sufficiently away from 0. In other words, $\frac{h(x)^2}{g(x)}$ converges as either $x \rightarrow 0$ or $\rightarrow \infty$. Being continuous, $\frac{h(x)^2}{g(x)}$ is upper bounded, and we obtain the desired inequality directly.

B. Supplementary Simulation Results on Different Graph Topologies

We conducted supplementary simulation experiments for the three comparison graph topologies studied by simulation in [Shah et al. \(2016\)](#), which are listed below. Unfortunately, there does not exist a parameter modulating the connectivity for the three graph topologies, such as the comparison width W in the banded comparison graphs. We instead evaluated via simulations the ℓ_2 error bound of the maximum likelihood BTL parameter estimator as a function of the dynamic range parameter B and the number of items d and compared to the rate of increase associated to the theoretical bounds from Theorem 4.1 and Theorem 2, [Shah et al. \(2016\)](#). Similarly to the simulation settings in [Shah et al. \(2016\)](#), the true score parameters w^* were obtained as a normalized sample from the standard d -dimensional Gaussian distribution: for each $i = 1, \dots, d$, we sampled $w'_i \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(0, 1)$ and then set $w_i^* = B \cdot w'_i / \|w'\|_\infty$ to ensure a dynamic range value of B . The three kinds of comparison graphs and simulation settings for this study are as follow:

1. *Complete graph*: every pair of d items are compared $T = 5$ times. B ranges from e^{-1} to 1, and d from $\lfloor 100e^{-1} \rfloor$ to 100.




Figure 3. Illustrations of the three simulated comparison graphs with $d = 8$. (a) complete graph, (b) star graph, and (c) barbell graph.




Figure 4. ℓ_2 error versus the dynamic range parameter B . Average empirical ℓ_2 estimation error evaluated via simulations for (a) complete graphs, (b) star graphs, and (c) barbell graphs. The value of B ranges from e^{-1} to 1 for all three graphs, where d was set to 100 for the complete graph and to 20 for the other graphs. The figures also depict the rate of increase of the theoretical upper bounds based on Theorem 4.1 and Theorem 2 of Shah et al. (2016).

2. *Star graph*: one *hub* item is compared $T = 200$ times against every other item, and there is no comparison among the $d - 1$ non-hub items. B ranges from e^{-1} to 1, and d from $\lfloor 20e^{-1} \rfloor$ to 20.
3. *Barbell graph*: for even d , every pair of items $1, \dots, d/2$ are compared $T = 100$ times, and so are items $d/2 + 1, \dots, d$. Items $d/2$ and $d/2 + 1$ are also compared $T = 100$ times. B ranges from e^{-1} to 1, and d from $\lfloor 20e^{-1} \rfloor$ to 20.

See Fig. 3 for the illustrations of the simulated comparison graphs when $d = 8$. For each of the three graph topologies and a grid of values for B and d , we simulated 100 datasets, each time with a different draw for w^* . The maximum likelihood BTL parameter estimate for each dataset was found as described in Section 5. Since a different value of w^* was used for each dataset, a direct comparison between the average estimated ℓ_2 errors and the theoretical error bounds can be misleading. Instead, we evaluate the 95th percentile of the empirical ℓ_2 errors over the simulated datasets for each combination of (B, d) . Fig. 4 and Fig. 5 show the 95th percentile curves of the empirical ℓ_2 errors for the three graph typologies shown in Fig. 3 as a function of the dynamic range B and of the number of items d , respectively. On each plot we also depict the theoretical bounds from Theorem 4.1 and Theorem 2 of Shah et al. (2016) arbitrarily shifted for better readability. Thus, the curves are to be compared based on their rates of growth of and not their values. The plot demonstrate that, for a fixed d and varying B , the error rates predicted by Theorem 4.1 track closely the estimated ℓ_2 errors and are more accurate than the ones implied by Theorem 2 of Shah et al. (2016) in all the graph topologies under consideration. On the other hand, the two error bounds appear to be similar and possibly loose as a function of d for a fixed value of B . Fig. 4 shows

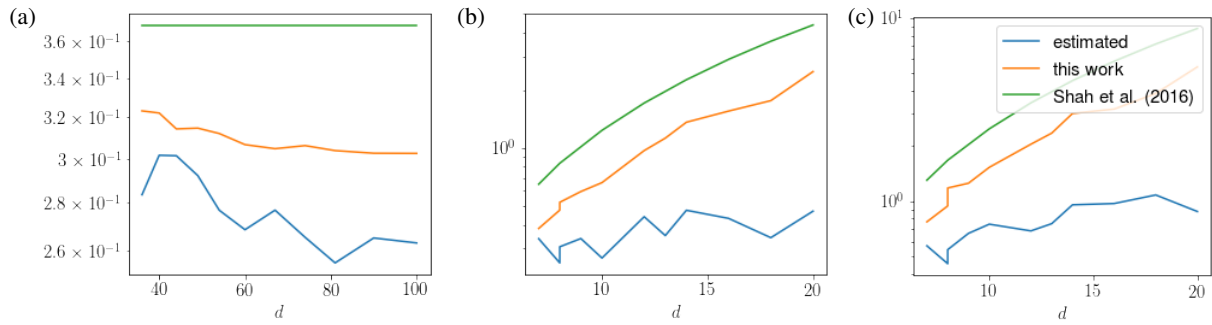


Figure 5. ℓ_2 error versus the number of items d . Average empirical ℓ_2 estimation error evaluated via simulations for (a) complete graphs, (b) star graphs, and (c) barbell graphs. For complete graphs, d ranges from $\lfloor 100e^{-1} \rfloor$ to 100, where the range is from $\lfloor 20e^{-1} \rfloor$ to 20 for the others. B was fixed at 1 for every case. The figures also depict the rate of increase of the theoretical upper bounds based on Theorem 4.1 and Theorem 2 of Shah et al. (2016).