Paper Review:
Rank Centrality: Ranking from Pair-wise Comparisons

Narutatsu (Edward) Ri

1 Problem Setup

The paper addresses the rank aggregation problem of determining the rankings between n items and their weights from observing a finite (and hence, noisy) number of pairwise comparison results. Here, the learner observes k comparisons between selected item pairs (i, j)[1] that yield outcomes such as "j beats i." A common modeling assumption for the rank aggregation problem (and by this paper) is that the observed pairwise comparisons are sampled according to the *Bradley-Terry-Luce (BTL) model*—or, more generally, from the pairwise marginals of the Plackett–Luce model. See Appendix A for an overview of their relationship.

Bradley-Terry-Luce Model. Each item $i \in [n]$ has a true but unknown positive weight w_i that represents "how preferred" the item is. The probability that j is preferred over i in a comparison is

$$\Pr[j \succ i] = \frac{w_j}{w_i + w_j}.$$

(For simplicity, we use $i \succ j$ to denote "i beats j" for the rest of the scribe.)

Observed outcomes are conditionally independent given these unknown true scores $\{w_i\}$. The scale of the weights are also unimportant, and we can assume $\sum_{i\in[n]} w_i = 1$ without loss of generality. See Appendix B for other (less common) models for ranking.

2 Existing Methods

2.1 Straightforward Approaches

Borda Count. A simple but naive approach would involve simply estimating the average ratio an item wins against others and using this as the weight. This is what Borda counting does. Concretely, let

$$a_{ij} = \begin{cases} \text{the fraction of times } i \text{ beats } j, & \text{if } i \text{ and } j \text{ were compared;} \\ \text{undefined or zero,} & \text{if } i \text{ and } j \text{ have not been compared.} \end{cases}$$

Note that $a_{ij} + a_{ji} = 1$ whenever i and j are compared. The Borda count assigns a Borda score to item i by the ratio of its wins, i.e.,

$$\text{Borda Score}(i) = \frac{\sum_{j\neq i} a_{ij}}{\sum_{j\neq i} \mathbb{1}\{(i,j) \text{ are compared}\}},$$

and the final ranking is obtained by sorting items in descending order of their Borda scores.

Borda count was historically used for complete rankings when sets of full orderings were available, where the i-th ranking item receives $n - i$ points and points are summed over all orderings to obtain Borda scores for each item. However, Borda count only considers "local" measures of rankings and hence a biased (both in the finite

[1] Note that in total, there are $\binom{n}{2} \cdot k$ observations that can be made. Here, the learner observes results for a subset of the full $\binom{n}{2}$ pairs.

and asymptotic regime) and inconsistent estimator.[2] For example, it turns out that in some cases, Borda count can suffer a constant error rate independent of k (the number of observed comparisons). See Appendix C for a constructive bad case analysis.

Maximum Likelihood Estimation (MLE). Another intuitive approach would be to compute the MLE with respect to $w_1, \ldots, w_n$ given the observed comparisons. Concretely, let Y_ℓ be a binary random variable for the ℓ-th comparison:

$$Y_\ell = \begin{cases} 1 & \text{if } j \succ i, \\ 0 & \text{otherwise.} \end{cases}$$

To denote which items are compared and in which order, we construct a design vector x_ℓ where $x_\ell = e_j - e_i$ if the ℓ-th comparison involves items i and j and e_i is the i-th standard basis vector. Under the BTL model, $\Pr[Y_\ell = 1] = w_j/(w_i + w_j), \Pr[Y_\ell = 0] = w_i/(w_i + w_j)$. Reparameterizing w_i as $\theta_i = \log w_i$:

$$\frac{w_j}{w_i + w_j} = \frac{e^{\theta_j}}{e^{\theta_i} + e^{\theta_j}} = \frac{1}{1 + e^{\theta_i - \theta_j}},$$

which is the form of a logistic function. Given the above, the probability for a single observation (ℓ-th comparison between items i and j) can be written as:

$$\begin{aligned} \Pr[Y_\ell \mid \boldsymbol{\theta}] &= \left[\frac{1}{1 + \exp(\theta_i - \theta_j)}\right]^{Y_\ell} \cdot \left[\frac{1}{1 + \exp(\theta_j - \theta_i)}\right]^{1 - Y_\ell} \\ &= \frac{\exp(Y_\ell \langle \boldsymbol{\theta}, x_\ell \rangle)}{1 + \exp(\langle \boldsymbol{\theta}, x_\ell \rangle)}, \end{aligned}$$

where the second line follows from using the inner product notation with $x_\ell = e_j - e_i$. Denoting the total number of (i.i.d.) comparisons as m, the overall log-likelihood function is given by:

$$L(\boldsymbol{\theta}) = \sum_{\ell=1}^{m} \left\{ Y_\ell \langle \boldsymbol{\theta}, x_\ell \rangle - \log\left[1 + \exp(\langle \boldsymbol{\theta}, x_\ell \rangle)\right] \right\},$$

and we are interested in $\boldsymbol{\theta}_{\mathrm{MLE}} = \arg\max_{\boldsymbol{\theta}} L(\boldsymbol{\theta})$. However, taking the derivative and setting it to zero results in a set of n nonlinear and coupled equations, and we must resort to numerical optimization-based approaches (e.g., following the gradient $\theta^{(t+1)} = \theta^{(t)} + \eta \nabla_\theta L\big(\theta^{(t)}\big)$ for some step size η).[3]

2.2 Spectral Ranking Approaches

There exists a long line of spectral-based approaches for assigning scores from a set of pairwise preferences (Seeley, 1949; Kendall, 1955). We introduce the most notable one below.

MC1∼MC4. Dwork et al. (2001) proposed several Markov chain based spectral ranking algorithms that all follow a similar idea with different ways to construct the transition probabilities of a random walk on items and use its stationary distribution to rank them.

For example, for MC2, the transition kernel $\mathbf{P}$ is constructed as:

$$\mathbf{P}_{ij}^{(\mathrm{MC2})} = \frac{a_{ij}}{\sum_{\ell \neq i} a_{i\ell}},$$

[2] Another quick and dirty intuition is as follows: each win is treated equally in Borda count, whereas wins should be weighted according to the "strength" of the participating items (i.e., winning against an item with a low weight is not equivalent to winning against an item with a high weight).Hence, If item i has mostly competed against weak opponents, it can artificially get a high Borda score.

[3] Quite irrelevant, but note that in nonlinear estimation (especially in logistic regression), MLE is typically biased in the finite-sample regime (due to nonlinearity of the log-likelihood function and the use of a normalization constraint) but unbiased in the infinite-sample regime, and consistent.

where a_{ij} is the winrate ratio of $j \succ i$. Here, one normalizes the i-th row locally by its sum. In the limit, since $a_{ij} \to w_j/(w_i + w_j)$, the denominator becomes

$$\sum_{\ell \neq i} a_{i\ell} \to \sum_{\ell \neq i} \frac{w_\ell}{w_i + w_\ell},$$

which is not necessarily proportional to w_i or uniform across i. As a result, the stationary distribution of MC2 is biased and does not automatically satisfy $v_i \propto w_i$ unless additional structure is present.

MC3 and MC4 are, in essence, MC2 with a slightly different construction of the transition matrix. Namely,

$$\mathbf{P}_{ij}^{(\mathrm{MC3})} = \begin{cases} \frac{a_{ij}}{\deg(i)}, & i \neq j, \\ 1 - \frac{1}{\deg(i)} \sum_{\ell \neq i} a_{i\ell}, & i = j, \end{cases}$$

$$\mathbf{P}_{ij}^{(\mathrm{MC4})} = \begin{cases} \frac{1}{n}, & \text{if } a_{ij} \geq a_{ji}, \\ 0, & \text{if } a_{ij} < a_{ji}, \\ 1 - (\text{appropriate sum}), & \text{if } i = j. \end{cases}$$

where $\deg(i)$ is the number of items compared with i.

In all of these, the normalization is done locally (depending on the number of comparisons or the row-sum of a_{ij}'s). Although these algorithms are also spectral since their scores are given by the stationary distribution of $\mathbf{P}^{(\mathrm{MCx})}$, they do not necessarily satisfy the reversibility property that would guarantee that the true BTL scores are recovered in the limit.

3 Rank Centrality

Rank Centrality proceeds by first forming a comparison graph $G = ([n], E)$ in which an edge (i, j) exists if the pair (i, j) was compared. From the observed data, let

$$a_{ij} = \frac{1}{k} \sum_{\ell=1}^{k} Y_{ij}^{\ell},$$

be the empirical fraction that item j is preferred to item i (note that $a_{ij} + a_{ji} = 1$). Then, RC defines a weight on the directed edge from i to j by

$$A_{ij} = \frac{a_{ij}}{a_{ij} + a_{ji}},$$

so that $A_{ij} \to w_j/(w_i + w_j)$ in the infinite-sample limit. Then, RC constructs a Markov chain on the items with transition matrix $\mathbf{P}$ given by

$$\mathbf{P}_{ij}^{(\mathrm{RC})} = \begin{cases} \frac{1}{d_{\max}} A_{ij}, & \text{if } i \neq j, \\ 1 - \frac{1}{d_{\max}} \sum_{k \neq i} A_{ik}, & \text{if } i = j, \end{cases}$$

where $d_{\max}$ is the maximum out-degree in the comparison graph (note that $d_{\max} \geq \sum_j A_{ij}$ for all $i \in [n]$ by construction). $\mathbf{P}^{(\mathrm{RC})}$ turns out to be ergodic and hence has a unique stationary distribution, which can be simply computed via the power method:

$$p_{t+1}^T = p_t^T \, \mathbf{P}^{(\mathrm{RC})},$$

and $\lim_{t \to \infty} p_t = \boldsymbol{\pi}$. This is the final estimate of the weights of Rank Centrality.

Key Takeaway of RC. We emphasize that this choice of construction results in $\mathbf{P}$ to be reversible, which is what gives Rank Centrality its edge over MC1∼MC4 (and allows for a lot of nice spectral-based analysis—see, e.g., Lemma 2). Concretely, if we substitute the ideal limits $A_{ij} = \frac{w_j}{w_i+w_j}$ into $\mathbf{P}_{ij}^{(\mathrm{RC})}$, then for $i \neq j$,

$$\widetilde{\mathbf{P}}_{ij}^{(\mathrm{RC})} = \frac{1}{d_{\max}} \frac{w_j}{w_i + w_j}.$$

Denoting the stationary distribution as $\widetilde{\boldsymbol{\pi}}$ (we can show that the resulting Markov Chain is ergodic), we obtain

$$\widetilde{\boldsymbol{\pi}}(i) = \frac{w_i}{\sum_j w_j},$$

which is exactly the target weights to estimate.

3.1 Properties of Rank Centrality

Some "fun facts:" Rank Centrality (RC) is biased in the finite-sample regime but becomes asymptotically unbiased. Empirically, RC also outperforms other spectral ranking approaches when data are randomly missing, and turns out to be an order of magnitude faster than MLE (and hence more preferred). Below, we present some theoretical guarantees of RC in full detail.

3.1.1 General Graphs

Theorem 1. *Denote* $b = \max_{i,j} \frac{w_i}{w_j}$, *and* $\kappa = \frac{d_{\max}}{d_{\min}}$. *For* $k \geq \left(\frac{256\, b^5 \kappa^2}{d_{\max} \xi^2}\right) \log n$, *with probability at least* $1 - 4/n$, *the normalized error is bounded from above by:*

$$\frac{\left\|\boldsymbol{\pi} - \widetilde{\boldsymbol{\pi}}\right\|}{\|\widetilde{\boldsymbol{\pi}}\|} \leq \frac{8\, b^{5/2} \kappa}{\xi} \cdot \sqrt{\frac{\log n}{k\, d_{\max}}}.$$

We first try to gain an intuitive understanding by considering the role each variable plays.

- b: Measures the spread among the true weights. Intuitively, the larger b is, the more "skewed" the MC is (i.e., there are states that are visited significantly more frequently than others)
- κ: Quantifies the imbalance in how many comparisons different items have
- ξ: Spectral gap. Measures the connectedness of the graph. Importantly, this is the spectral gap of the graph with equal probability weightings for each outgoing edge, and not $\widetilde{\mathbf{P}}$.

It may also be worth noting that some of these parameters are not independent of each other.[4]

The intuition for the theorem is as follows: We are interested in the deviation of $\boldsymbol{\pi}$ from $\widetilde{\boldsymbol{\pi}}$, i.e., when we traverse a MC with transition kernel $\mathbf{P}^{(\mathrm{RC})}$ sufficiently, how much further off is the observed distribution from the "true" distribution. Intuitively, this will be largely dependent on the "deviation" of $\mathbf{P}^{(\mathrm{RC})}$ from the true value $\widetilde{\mathbf{P}}^{(\mathrm{RC})}$, which is denoted as Δ. If Δ is "small enough" by some notion, then it may be possible to shave off the deviation by sufficiently traversing the MC. This is precisely what Theorem 1 is formalizing.

Proof Sketch. Proving Theorem 1 is done via three steps:

1. Formalizing how the distribution edges closer to $\widetilde{\boldsymbol{\pi}}$ as $t \to \infty$ (Lemma 2). We expect this to largely depend on Δ. It turns out that this also involves a variable ρ, that depends on Δ and the spectral gap ξ.
2. Bounding Δ (Lemma 3).
3. Bounding ρ (Lemma 4).

[4]Note that in the original paper, there is a constant C which is set to be $C \geq 8$ and appears in the bound. However, it is not very important, so we will replace C with 8 in this scribe for some parts.

Lemma 2. *Consider a MC with transition kernel $\mathbf{P} = \widetilde{\mathbf{P}} + \Delta$ where $\widetilde{\mathbf{P}}$ is reversible, and let p_t be the distribution of the MC $\mathbf{P}$ when started with initial distribution p_0. Then,*

$$\frac{\left\|p_t - \widetilde{\boldsymbol{\pi}}\right\|}{\|\widetilde{\boldsymbol{\pi}}\|} \leq \underbrace{\rho^t \frac{\|p_0 - \widetilde{\boldsymbol{\pi}}\|}{\|\widetilde{\boldsymbol{\pi}}\|} \sqrt{\frac{\widetilde{\boldsymbol{\pi}}_{\max}}{\widetilde{\boldsymbol{\pi}}_{\min}}}}_{\text{mixing towards ideal chain}} + \underbrace{\frac{1}{1-\rho} \|\Delta\|_2 \sqrt{\frac{\widetilde{\boldsymbol{\pi}}_{\max}}{\widetilde{\boldsymbol{\pi}}_{\min}}}}_{\text{unavoidable fixed error floor}} .$$

where $\widetilde{\boldsymbol{\pi}}_{\min} = \min_i \widetilde{\boldsymbol{\pi}}(i)$, $\widetilde{\boldsymbol{\pi}}_{\max} = \max_i \widetilde{\boldsymbol{\pi}}(i)$, and $\rho = \lambda_{\max}(\widetilde{\mathbf{P}}) + \|\Delta\|_2 \sqrt{\frac{\widetilde{\boldsymbol{\pi}}_{\max}}{\widetilde{\boldsymbol{\pi}}_{\min}}}$.

As seen above, the error term splits into two parts. The first term will tend towards 0 (given $\rho < 1$, which turns out to be true given the conditions on k), so we only need to focus on bounding the second term. This turns out to be non-trivial, and requires individually bounding $\|\Delta\|_2$ and $1 - \rho$ in Lemmas 3 and 4 respectively.

Proof Sketch. $\widetilde{\mathbf{P}}$ is reversible and satisfies the detailed balance condition $\widetilde{\boldsymbol{\pi}}(i)\, \widetilde{\mathbf{P}}_{ij} = \widetilde{\boldsymbol{\pi}}(j)\, \widetilde{\mathbf{P}}_{ji}$ for all i, j. This reversibility allows us to introduce the weighted inner product space

$$L^2(\widetilde{\boldsymbol{\pi}}) = \left\{ a \in \mathbb{R}^n \ : \ \langle a, b \rangle_{\widetilde{\boldsymbol{\pi}}} = \sum_{i=1}^{n} a_i\, \widetilde{\boldsymbol{\pi}}(i)\, b_i \right\}$$

in which the corresponding norm is defined by $\|a\|_{\widetilde{\boldsymbol{\pi}}} = \sqrt{\sum_{i=1}^n \widetilde{\boldsymbol{\pi}}(i)\, a_i^2}$. Because the inner product weights each coordinate by $\widetilde{\boldsymbol{\pi}}(i)$, working in this space naturally "respects" the stationary measure. The operator norm in this space is defined as:

$$\|A\|_{\widetilde{\boldsymbol{\pi}},2} = \max_{a \neq 0} \frac{\|Aa\|_{\widetilde{\boldsymbol{\pi}}}}{\|a\|_{\widetilde{\boldsymbol{\pi}}}} .$$

Also note that for any $a \in \mathbb{R}^n$ one has

$$\begin{aligned} \widetilde{\boldsymbol{\pi}}_{\min}^{\frac{1}{2}} \|a\| \ \leq \ \|a\|_{\widetilde{\boldsymbol{\pi}}} \ \leq \ \widetilde{\boldsymbol{\pi}}_{\max}^{\frac{1}{2}} \|a\|, \\ \widetilde{\boldsymbol{\pi}}_{\min} \|A\|_2 \ \leq \ \|A\|_{\widetilde{\boldsymbol{\pi}},2} \ \leq \ \widetilde{\boldsymbol{\pi}}_{\max} \|A\|_2 . \end{aligned}$$

These inequalities are used to convert back to the Euclidean norm later on in the proof.

Next, define the symmetrized matrix using the diagonal matrix $\widetilde{\boldsymbol{\pi}}$ with entries $\widetilde{\boldsymbol{\pi}}_{ii} = \widetilde{\boldsymbol{\pi}}(i)$ as $\mathbf{S} = \widetilde{\boldsymbol{\pi}}^{\frac{1}{2}}\, \widetilde{\mathbf{P}}\, \widetilde{\boldsymbol{\pi}}^{-\frac{1}{2}}$. Because $\widetilde{\mathbf{P}}$ is reversible, one can verify that $\mathbf{S}$ is symmetric, which implies that $\mathbf{S}$ has real eigenvalues with largest eigenvalue of 1 with multiplicity one and the remaining eigenvalues lie in the interval $[-1, 1]$. Denote the eigenvalues of $\mathbf{S}$ by

$$1 = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n \geq -1,$$

and define $\lambda_{\max} = \max\{|\lambda_n|, \lambda_2\}$. The key idea in this proof is to use the spectral properties of $\mathbf{S}$ (and hence of $\widetilde{\mathbf{P}}$) to show that after "projecting out" the dominant eigencomponent (which corresponds to the stationary distribution) the remainder of the operator contracts.

Namely, let u_1 be the eigenvector of $\mathbf{S}$ corresponding to $\lambda_1 = 1$. Because of the relationship between $\mathbf{S}$ and $\widetilde{\mathbf{P}}$, the corresponding left eigenvector of $\widetilde{\mathbf{P}}$ is $v_1 = \widetilde{\boldsymbol{\pi}}^{\frac{1}{2}}\, u_1$, which is equal to $\widetilde{\boldsymbol{\pi}}$ (up to normalization) since we know that $\widetilde{\boldsymbol{\pi}}$ is the stationary distribution. One may also note that $\widetilde{\boldsymbol{\pi}}^{-\frac{1}{2}}\, u_1 = \mathbf{1}$. Then the projection (or "rank-one" operator) onto the dominant (stationary) component is defined by

$$\mathbf{S}_1 = u_1\, u_1^T,$$

and accordingly,

$$\widetilde{\mathbf{P}}_1 = \widetilde{\boldsymbol{\pi}}^{-\frac{1}{2}}\, \mathbf{S}_1\, \widetilde{\boldsymbol{\pi}}^{\frac{1}{2}} .$$

Because $\widetilde{\mathbf{P}}_1$ captures only the stationary distribution, the error $p_t - \widetilde{\boldsymbol{\pi}}$ is orthogonal to this rank-one space. In other words, if we start with any probability distribution p_0, then

$$(p_0 - \widetilde{\boldsymbol{\pi}})^T\, \widetilde{\boldsymbol{\pi}}^{-\frac{1}{2}}\, u_1 = (p_0 - \widetilde{\boldsymbol{\pi}})^T\, \mathbf{1} = 0,$$

so that

$$(p_0 - \widetilde{\boldsymbol{\pi}})^T \, \widetilde{\mathbf{P}}_1 \;=\; 0.$$

Now, consider the actual noisy chain $\mathbf{P}$. Noting that $p_t^T = p_0^T \, \mathbf{P}^t$, a first-order recursion for the difference between p_t and $\widetilde{\boldsymbol{\pi}}$ can be written as:

$$\begin{aligned} p_t^T - \widetilde{\boldsymbol{\pi}}^T &= (p_{t-1} - \widetilde{\boldsymbol{\pi}})^T \, \mathbf{P} + \widetilde{\boldsymbol{\pi}}^T \, \Delta \\ &= (p_{t-1} - \widetilde{\boldsymbol{\pi}})^T \, (\widetilde{\mathbf{P}} + \Delta) + \widetilde{\boldsymbol{\pi}}^T \, \Delta \\ &= (p_{t-1} - \widetilde{\boldsymbol{\pi}})^T \, (\widetilde{\mathbf{P}} - \widetilde{\mathbf{P}}_1 + \Delta) + \widetilde{\boldsymbol{\pi}}^T \, \Delta. \qquad (\because \forall p_\ell : (p_\ell - \widetilde{\boldsymbol{\pi}})^T \, \widetilde{\mathbf{P}}_1 \;=\; 0) \end{aligned}$$

The term $\widetilde{\mathbf{P}} - \widetilde{\mathbf{P}}_1$ represents the part of $\widetilde{\mathbf{P}}$ acting on the orthogonal complement of the stationary vector with operator norm exactly $\lambda_{\max}$. Thus, we obtain an inequality in the $L^2(\widetilde{\boldsymbol{\pi}})$ norm:

$$\|p_t - \widetilde{\boldsymbol{\pi}}\|_{\widetilde{\boldsymbol{\pi}}} \;\le\; \|p_{t-1} - \widetilde{\boldsymbol{\pi}}\|_{\widetilde{\boldsymbol{\pi}}} \left(\|\widetilde{\mathbf{P}} - \widetilde{\mathbf{P}}_1\|_{\widetilde{\boldsymbol{\pi}},2} + \|\Delta\|_{\widetilde{\boldsymbol{\pi}},2} \right) + \|\widetilde{\boldsymbol{\pi}}^T \, \Delta\|_{\widetilde{\boldsymbol{\pi}}}.$$

Defining $\rho = \lambda_{\max} + \|\Delta\|_{\widetilde{\boldsymbol{\pi}},2}$, the first term contracts by a factor ρ per iteration. Iterating this recursion yields

$$\|p_t - \widetilde{\boldsymbol{\pi}}\|_{\widetilde{\boldsymbol{\pi}}} \;\le\; \rho^t \, \|p_0 - \widetilde{\boldsymbol{\pi}}\|_{\widetilde{\boldsymbol{\pi}}} + \sum_{\ell=0}^{t-1} \rho^{t-1-\ell} \, \|\widetilde{\boldsymbol{\pi}}^T \, \Delta\|_{\widetilde{\boldsymbol{\pi}}}.$$

We then convert back to Euclidean space:

$$\|p_t - \widetilde{\boldsymbol{\pi}}\| \;\le\; \frac{\widetilde{\boldsymbol{\pi}}_{\max}^{\frac{1}{2}}}{\widetilde{\boldsymbol{\pi}}_{\min}^{\frac{1}{2}}} \, \rho^t \, \|p_0 - \widetilde{\boldsymbol{\pi}}\| \;+\; \frac{1}{1-\rho} \, \frac{\widetilde{\boldsymbol{\pi}}_{\max}^{\frac{1}{2}}}{\widetilde{\boldsymbol{\pi}}_{\min}^{\frac{1}{2}}} \, \|\Delta\|_2,$$

which provides the desired result. □

Lemma 3. *The error matrix $\Delta = P - \widetilde{\mathbf{P}}$ satisfies $\|\Delta\|_2 \le C\sqrt{\frac{\log n}{k \, d_{\max}}}$ with probability at least $1 - 4/n$.*

Proof Sketch. First, they break Δ down into $\Delta = D + \bar{\Delta}$, where $D = \mathrm{diag}(\Delta), \bar{\Delta} = \Delta - D$. Note that $\|\Delta\| \le \|D\| + \|\bar{\Delta}\|$ by triangle inequality. Now,

$$\Delta_{ij} = \begin{cases} \frac{1}{k \, d_{\max}} C_{ij} & \text{if } i \neq j \\ -\frac{1}{k \, d_{\max}} \sum_{j \neq i} C_{ij} & \text{otherwise,} \end{cases}$$

where $C_{ij} \sim \mathrm{Binom}(k, p_{ij}) - k p_{ij}$ (notice the k in the denominator of the expressions to counteract). The value for Δ_{ii} follows from

$$\Delta_{ii} = P_{ii} - \mathbf{P}_{ii} = \left(1 - \sum_{j \neq i} P_{ij}\right) - \left(1 - \sum_{j \neq i} \mathbf{P}_{ij}\right) = \sum_{j \neq i} \mathbf{P}_{ij} - \sum_{j \neq i} P_{ij} = -\sum_{j \neq i} \Delta_{ij} \, .$$

Note that $C_{ij} + C_{ji} = 0$ as each "error" in one direction is counterbalanced by an error in the reverse direction, and each C_{ij} has mean zero ($\mathbb{E}\big[C_{ij}\big] = k \, p_{ij} - k \, p_{ij} = 0$).

Bounding $\|D\|$. For fixed i, the random variable

$$S_i := k \, d_{\max} \, \Delta_{ii} = -\sum_{j \neq i} C_{ij}$$

is the sum of at most $k \, d_{\max}$ independent zero-mean random variables. Now, we can view the sequential partial sums of C_{ij} as forming a martingale. Suppose we order the neighbors of i arbitrarily as $j_1, j_2, \ldots, j_N$, where $N \le k \, d_{\max}$. Define the partial sums

$$S_0 = 0, \quad S_\ell = -\sum_{r=1}^{\ell} C_{ij_r}$$

for $\ell = 1, \ldots, N$, where C_{ij_r} represents the r-th comparison (i.i.d. draw) out of the N total comparisons (draws) from $\text{Bern}(p_{ij}) - p_{ij}$. Because each C_{ij_r} has mean zero, we have for each ℓ,

$$\mathbb{E}\Big[S_\ell \,\Big|\, C_{ij_1}, \ldots, C_{ij_{\ell-1}}\Big] = S_{\ell-1},$$

which is the very definition of a martingale. Thus, we apply Azuma-Hoeffding, which tells us that for any $t > 0$,

$$\begin{aligned} &\Pr\left[|S_i| > t\right] \le 2\exp\left(-\frac{t^2}{2\,k\,d_{\max}}\right) \\ \iff &\Pr\left[|\Delta_{ii}| > \frac{t}{k\,d_{\max}}\right] \le 2\exp\left(-\frac{t^2}{2\,k\,d_{\max}}\right). \end{aligned}$$

Setting $t \ge C\,k\,d_{\max}\,\log n$ for some constant C, we obtain:

$$\Pr\left[|\Delta_{ii}| > C\,\frac{\log n}{k\,d_{\max}}\right] \le 2\exp\left(-\frac{C^2\,k\,d_{\max}}{2}\log n\right) \le 2\,n^{-C^2/2}.$$

Taking a union bound over all i (there are n indices), we conclude:

$$\Pr\left[\|D\|_2 \ge C\,\frac{\log n}{k\,d_{\max}}\right] \le 2\,n\cdot n^{-C^2/2} = 2\,n^{-C^2/2+1}.$$

Thus, with high probability, the operator norm of the diagonal part D is bounded by the desired term.

Bounding $\|\bar{\Delta}\|$. Here, they consider the two cases $d_{\max} \le \log n$ and otherwise separately. This is because when $d_{\max} \le \log n$, each row of the matrix involves a modest number of random variables, and a direct bound on the row sums (e.g., Azuma–Hoeffding + union bound) produces a tight overall bound. However, when each row sums many terms, taking a union bound over many items can become very loose, and a finer and more matrix-aware would provide a tigher bound.

For the case where $d_{\max} \le \log n$, they invoke the inequality $\|M\|_2 \le \sqrt{\|M\|_1\|M\|_\infty}$ for any square matrix M, where the two terms inside represent the maximal row- and column-sum of absolute values of M. Hence, define the variable

$$R_i = \frac{1}{k\,d_{\max}}\sum_{j\neq i}|C_{ij}|,$$

which measures the total absolute error in the ith row, so we wish to bound $\Pr[R_i > s]$ for arbitrary s. Rewriting,

$$\Pr\left[R_i > s\right] = \Pr\left[\sum_{j\neq i}|C_{ij}| > k\,d_{\max}\,s\right] \le \sum_{j\in\partial i}\sum_{\xi_j\in\{-1,+1\}}\Pr\left[\sum_{j\in\partial i}\xi_j\,C_{ij} > k\,d_{\max}\,s\right],$$

where the second line follows from $Pr\left[|X| > a\right] \le \Pr\left[X > a\right] + \Pr\left[-X > a\right]$, which comes from $|x| \le \max\{x, -x\}$. Here, the sum over ξ_j (each taking values -1 and $+1$) is a union bound over the two "sides" for each j, of which there are $2d_i$ such terms. Since $d_i \le d_{\max}$, we can bound this total factor by $\exp\big(d_{\max}\ln 2\big)$.

Noting that C_{ij} is the sum of k independent bounded differences, we apply Hoeffding's inequality:

$$\begin{aligned} \Pr\left[\sum_{j\in\partial i}\xi_j\,C_{ij} > k\,d_{\max}\,s\right] &\le \exp\left(-\frac{(k\,d_{\max}\,s)^2}{2\,d_{\max}}\right) && (\because d_i \le d_{\max}\text{ terms}) \\ &\le \exp\big(-2\,k\,d_{\max}\,s^2\big), \end{aligned}$$

Thus,

$$\Pr\left[R_i > s\right] \le 2d_i\,\exp\big(-2\,k\,d_{\max}\,s^2\big) \le \exp\big(-2\,k\,d_{\max}\,s^2 + d_{\max}\ln 2\big).$$

Then, applying a union bound over all n rows, we get

$$\Pr\left[\|\overline{\Delta}\|_2 \geq s\right] \leq 2n\, \exp\left(-2\, k\, d_{\max}\, s^2 + d_{\max} \ln 2\right).$$

Choosing $s = \frac{C_2 \log n + d_{\max} \ln 2}{k\, d_{\max}}$ and plugging this value into the bound shows that the tail probability is small (on the order of $n^{-(C_2/2-1)}$), and hence with high probability we have

$$\|\overline{\Delta}\|_2 \leq \frac{C_2 \log n}{k\, d_{\max}},$$

where we also used the assumption that

$$d_{\max} \leq \log n,$$

so that $\log n + d_{\max} \ln 2$ can be thought of as $\Theta(\log n)$.

For $d_{\max} > \log n$, a result for self-adjoint matrix martingales is applied. First, they rewrite $\overline{\Delta} = \sum_{i<j} Z_{ij}$, where each Z_{ij} is the contribution from pair (i, j) and is given by:

$$Z_{ij} = \left(e_i\, e_j^T - e_j\, e_i^T\right)\left(P_{ij} - \mathbf{P}_{ij}\right).$$

Since the Z_{ij} are independent over unordered pairs (with $i < j$), they form a sum of random matrices.

Z_{ij} is generally not self-adjoint because they are skew-symmetric. To apply the matrix Bernstein inequality which requires the summands to be self-adjoint, they "dilate" each Z_{ij} into a block matrix of the form:

$$\widetilde{Z}_{ij} = \begin{pmatrix} 0 & Z_{ij} \\ Z_{ij}^T & 0 \end{pmatrix},$$

which is self-adjoint and preserves the spectral norm ($\|Z_{ij}\|_2 = \|\widetilde{Z}_{ij}\|_2$). One must then show that these dilated matrices satisfy a moment condition of the form

$$\mathbb{E}\Big[(\widetilde{Z}_{ij})^p\Big] \preceq \frac{p!}{2}\, C^{p-2}\, (\widetilde{A}_{ij})^2,$$

for some constant C and auxiliary matrices $\widetilde{A}_{ij}$. Here, the fact that each Δ_{ij} scales as $1/(k\, d_{\max})$ and that each C_{ij} is sub-Gaussian by Hoeffding's inequality is used to prove such a bound. They also use a fact about the matrices A_{ij} that their powers are bounded because the only eigenvalues are ± 1.

Given that these constraints are satisfied, they apply the matrix Bernstein inequality whichgives a tail bound for $\left\|\sum_{i<j} \widetilde{Z}_{ij}\right\|_2$, which in turn implies

$$\|\overline{\Delta}\|_2 \leq \frac{C\, \log n}{k\, d_{\max}},$$

with high probability.[5] □

Lemma 4. *For the abovementioned conditions on k and $\|\Delta\|_2$, $1 - \rho \geq \frac{\xi d_{\min}}{b^2 d_{\max}}$ holds.*

Proof Sketch. Given that $\rho = \lambda_{\max}(\widetilde{\mathbf{P}}) + \|\Delta\|_2 \sqrt{\frac{\widetilde{\boldsymbol{\pi}}_{\max}}{\widetilde{\boldsymbol{\pi}}_{\min}}}$, plugging in the bound for $\|\Delta\|_2$, we get

$$1 - \rho = 1 - \lambda_{\max}(\widetilde{\mathbf{P}}) - \sqrt{\frac{C^2\, b \log n}{k d_{\max}}},$$

where we note that $\frac{\widetilde{\boldsymbol{\pi}}_{\max}}{\widetilde{\boldsymbol{\pi}}_{\min}} = \max_{i,j} \frac{w_i}{w_j} = b$.

[5] Note that the proof for the $d_{\max} > \log n$ setting can also be directly applied for the $d_{\max} \leq \log n$ case. It's unclear to me why they decided to separate the two cases aside from presentation purposes.

First, note that plugging in the expression for $k = \left(\frac{4\,C^2\,b^5\kappa^2}{d_{\max}\xi^2}\right)\log n$ into $\sqrt{\frac{C^2\,b\log n}{k d_{\max}}}$ gives $\frac{1}{2}\frac{\xi}{b^2}\frac{d_{\min}}{d_{\max}}$. Hence, we mainly focus on bounding $1-\lambda_{\max}(\widetilde{\mathbf{P}})$, which is essentially the spectral gap of the graph induced by the true transition kernel $\widetilde{\mathbf{P}}$. However, directly obtaining the spectral gaph for $\widetilde{\mathbf{P}}$ is infeasible, so we resort to the *comparison lemma*, which involves introducing a simpler chain where the spectral gap is well-understood, then relate the spectral gaps of this simpler chain to $\widetilde{\mathbf{P}}$.

Namely, define a new transition matrix $\mathbf{Q}$ corresponding to the simple random walk on the comparison graph:

$$\mathbf{Q}_{ij} = \begin{cases} \frac{1}{d_i} & \text{if } i \text{ and } j \text{ were compared} \\ 0 & \text{otherwise,} \end{cases}$$

where d_i is the degree (number of items compared with) of node i. The stationary distribution of $\mathbf{Q}$ is known to be

$$\boldsymbol{\mu}(i) = \frac{d_i}{\sum_j d_j}\,.$$

Because the simple random walk is reversible, its spectral gap is given by $1-\lambda_{\max}(\mathbf{Q}) = \xi$.

Here, the comparison lemma stating that for two reversible Markov chains $\mathbf{Q}$ and $\widetilde{\mathbf{P}}$ defined on the same graph with stationary measures $\boldsymbol{\mu}$ and $\widetilde{\boldsymbol{\pi}}$,

$$\frac{1-\lambda_{\max}(\widetilde{\mathbf{P}})}{1-\lambda_{\max}(\mathbf{Q})} \geq \frac{\alpha}{\beta} \iff 1-\lambda_{\max}(\mathbf{P}) \geq \frac{\alpha}{\beta}\,\xi$$

where

$$\alpha = \min_{(i,j)\in E}\left\{\frac{\widetilde{\boldsymbol{\pi}}(i)\,\mathbf{P}_{ij}}{\boldsymbol{\mu}(i)\,\mathbf{Q}_{ij}}\right\},$$
$$\beta = \max_i\left\{\frac{\widetilde{\boldsymbol{\pi}}(i)}{\boldsymbol{\mu}(i)}\right\}.$$

Thus, if we can lower-bound the ratio α/β in terms of $d_{\min}$, $d_{\max}$, and b, then we get a lower bound for $1-\lambda_{\max}(\mathbf{P})$.

For β, recall:

$$\widetilde{\boldsymbol{\pi}}(i) = \frac{w_i}{\sum_j w_j}, \qquad \boldsymbol{\mu}(i) = \frac{d_i}{\sum_j d_j}\,.$$

The assumption $\max_{i,j}\frac{w_i}{w_j} \leq b$ implies that the weights are all within a constant factor of each other. Therefore, one can show that

$$\beta \leq \frac{b\,|E|}{n\,d_{\min}},$$

where $|E|$ is the total number of edges. (This bound comes from relating the averages $\sum_j w_j$ and $\sum_j d_j$; roughly speaking, since d_i capture the number of comparisons and the chain $\widetilde{\mathbf{P}}$ is built using the weights, one can relate the two.)

For α, observe that by definition, $\widetilde{\mathbf{P}}_{ij} = \frac{1}{d_{\max}}\frac{w_j}{w_i+w_j}$, and $\mathbf{Q}_{ij} = \frac{1}{d_i}$. Hence, their ratio becomes

$$\frac{\widetilde{\boldsymbol{\pi}}(i)\,\widetilde{\mathbf{P}}_{ij}}{\boldsymbol{\mu}(i)\,\mathbf{Q}_{ij}} = \frac{\frac{w_i}{\sum_j w_j}\cdot\frac{1}{d_{\max}}\frac{w_j}{w_i+w_j}}{\frac{d_i}{\sum_j d_j}\cdot\frac{1}{d_i}} = \frac{w_i}{\sum_j w_j}\cdot\frac{w_j}{w_i+w_j}\cdot\frac{\sum_j d_j}{d_{\max}}\,.$$

Using the fact that all w_i are within a factor of b and a lower bound for any positive numbers w_i and w_j that $\frac{w_i\,w_j}{w_i+w_j} \geq \frac{1}{2}\min\{w_i, w_j\}$, one can show that

$$\alpha \geq \frac{|E|}{n\,b\,d_{\max}}\,.$$

Then,

$$\frac{\alpha}{\beta} \geq \frac{d_{\min}}{b^2\,d_{\max}}\,.$$

Plugging back in the comparison lemma we have:

$$1 - \lambda_{\max}(\widetilde{\mathbf{P}}) \geq \frac{\alpha}{\beta}\,\xi \geq \frac{\xi\, d_{\min}}{b^2\, d_{\max}}\,.$$

Hence, $1 - \rho \geq \frac{1}{2}\frac{\xi\, d_{\min}}{b^2\, d_{\max}}$, but we absorb the constant and simply write

$$1 - \rho \geq \frac{\xi\, d_{\min}}{b^2\, d_{\max}}.$$

□

□

3.1.2 Random Graphs

Results for the (Erdos-Renyi) random graph setting trivially follows from Theorem 1. Namely, each pair (i, j) is selected to be in E with probability d/n independently where d denotes the average degree in the comparison graph, and each such chosen pair of objects is compared k times with the outcomes of comparisons.

Theorem 5. *For $d \geq 640 \log n$ and $k\,d \geq 8192\, b^5 \log n$, the following bound on the error rate holds with probability at least $1 - \frac{10}{n}$:*

$$\frac{\left\|\boldsymbol{\pi} - \widetilde{\boldsymbol{\pi}}\right\|}{\|\widetilde{\boldsymbol{\pi}}\|} \leq 64\, b^{5/2} \sqrt{\frac{\log n}{k\, d}}\,.$$

3.1.3 Optimality of Estimates

Information-Theoretic Lower Bounds (Minimax Framework). Here, they consider the setting where an algorithm $\mathcal{A}$ provides an estimate $\boldsymbol{\pi}^{\mathcal{A}}$ of the true weights $\widetilde{\boldsymbol{\pi}}$, and an adversary is allowed to choose such an $\widetilde{\boldsymbol{\pi}}$ where the (relative) error of $\mathcal{A}$ is maximixed. In particular, they consider the best error estimable by any algorithm $\mathcal{A}$ under this setting (i.e., minimax framework). This provides an idea of the best worst-case performance of any algorithm on the weight estimation task.

Theorem 6. *Let $\mathcal{S}_b$ denote the space of all BTL score vectors with dynamic range at most b:*

$$\mathcal{S}_b \equiv \left\{ \widetilde{\boldsymbol{\pi}} \in \mathbb{R}^n \;\middle|\; \sum_{i \in [n]} \widetilde{\boldsymbol{\pi}}_i = 1\,,\; \max_{i,j} \frac{\widetilde{\boldsymbol{\pi}}_i}{\widetilde{\boldsymbol{\pi}}_j} \leq b \right\},$$

for $b \geq 1$. Then,

$$\inf_{\mathcal{A}} \sup_{\widetilde{\boldsymbol{\pi}} \in \mathcal{S}_b} \frac{\mathbb{E}\big[\, \|\boldsymbol{\pi}^{\mathcal{A}} - \widetilde{\boldsymbol{\pi}}\| \,\big]}{\|\widetilde{\boldsymbol{\pi}}\|} \geq \frac{1}{240\sqrt{10}} \frac{b-1}{b+1} \frac{1}{\sqrt{kd}}\,,$$

under the Erdos-Renyi model.

Note that the bound has no dependency on n.

Proof Sketch. The first step is to construct a packing of vectors in S_b that are well separated. In other words, there exists a set $\{\, \widetilde{\boldsymbol{\pi}}(1), \widetilde{\boldsymbol{\pi}}(2), \ldots, \widetilde{\boldsymbol{\pi}}(M) \,\}$ such that any two distinct vectors satisfy

$$\|\widetilde{\boldsymbol{\pi}}(\ell_1) - \widetilde{\boldsymbol{\pi}}(\ell_2)\| \geq \delta,$$

for some separation parameter δ. It turns out that one can construct such a packing with $M = e^{n/128}$ elements (provided that $\delta \leq 1/(2^{10} n)$) up to constant factors.

Imagine that the adversary chooses one of these M vectors uniformly at random. Call this index $\widehat{L}$ and let the corresponding true vector be $\widetilde{\boldsymbol{\pi}}(\widehat{L})$. Given the BTL model, we obtain observations X (the outcomes of pairwise comparisons) that are generated according to the chosen true vector. Now, define an index estimate as

$$\widehat{L}_{\text{est}} = \arg\min_{\ell \in [M]} \|\boldsymbol{\pi} - \widetilde{\boldsymbol{\pi}}(\ell)\|,$$

i.e., the $\widetilde{\boldsymbol{\pi}}$ in the set of M items that is closest to the estimate $\boldsymbol{\pi}$ in the ℓ_2 norm sense. As the packing is δ-separated, if the estimator $\boldsymbol{\pi}$ is not extremely close to the true $\widetilde{\boldsymbol{\pi}}(\widehat{L})$, then $\|\boldsymbol{\pi} - \widetilde{\boldsymbol{\pi}}(\widehat{L})\| \geq \frac{\delta}{2}$. Thus, the risk of the estimator can be lower bounded by

$$\mathbb{E}\left[\|\boldsymbol{\pi} - \widetilde{\boldsymbol{\pi}}(\widehat{L})\|\right] \geq \frac{\delta}{2} \Pr\left[\widehat{L}_{\text{est}} \neq \widehat{L}\right].$$

Now we need to error bound $\Pr\left[\widehat{L}_{\text{est}} \neq \widehat{L}\right]$. We use Fano's inequality for this, which relates the probability of error in a multi-way hypothesis testing problem to the mutual information $I(\widehat{L}; X)$ between the true hypothesis (here, $\widehat{L}$) and X. Here, one obtains an inequality of the form

$$\Pr\left[\widehat{L}_{\text{est}} \neq \widehat{L}\right] \geq 1 - \frac{I(\widehat{L}; X) + \log 2}{\log M}.$$

Substituting this into the lower bound for the expected error gives

$$\mathbb{E}\left[\|\pi - \widetilde{\boldsymbol{\pi}}(\widehat{L})\|\right] \geq \frac{\delta}{2}\left(1 - \frac{I(\widehat{L}; X) + \log 2}{\log M}\right).$$

Thus, to lower bound the minimax risk, we need to upper bound $I(\widehat{L}; X)$ by relating it to the pairwise KL divergences between the distributions induced by different candidate vectors in the packing.

By the data processing inequality and Jensen's inequality, one obtains

$$I(\widehat{L}; X) \leq \frac{1}{M^2} \sum_{\ell_1, \ell_2 \in [M]} \mathrm{KL}\Big(P_{X|\widetilde{\boldsymbol{\pi}}(\ell_1)} \,\big\|\, P_{X|\widetilde{\boldsymbol{\pi}}(\ell_2)}\Big).$$

Next, using properties of the BTL model (where pairwise items (i, j) are compared k times), one can upper bound the KL divergence between the outcomes under two different BTL models (with weight vectors $\widetilde{\boldsymbol{\pi}}(\ell_1)$ and $\widetilde{\boldsymbol{\pi}}(\ell_2)$) by an expression proportional to

$$k\left[\Big(\widetilde{\boldsymbol{\pi}}_i(\ell_1) - \widetilde{\boldsymbol{\pi}}_i(\ell_2)\Big)^2 + \Big(\widetilde{\boldsymbol{\pi}}_j(\ell_1) - \widetilde{\boldsymbol{\pi}}_j(\ell_2)\Big)^2\right],$$

up to a multiplicative constant. Summing these contributions over all $\binom{n}{2}$ pairs weighted by the fact that any pair is included with probability d/n, we obtain

$$I(\widehat{L}; X) \leq 2\, n\, k\, d\, \delta^2.$$

Now, choose the packing separation

$$\delta \sim \frac{1}{30}\left(\frac{b+1}{b-1}\right)^{1/10}\Big(k\, d\, n\Big)^{-1/10},$$

with the precise constant chosen so that the term $I(\widehat{L}; X) + \log 2$ remains a fraction less than $\log M$ (ensuring that the probability of error is bounded away from zero). Since the size of the packing is $M = e^{n/128}$, we have $\log M \sim n/128$.

Substituting back into the Fano inequality, we get that every estimator must incur a normalized error at least on the order of $O\left(((b+1)/(b-1))^{1/10}\,(k\, d)^{-1}\right)$. □

Lower Bound via Cramer-Rao. We can also consider an "average-case" lower bound via a Cramer-Rao bound: the variance of any unbiased estimator $\hat{\theta}$ of θ is bounded by the reciprocal of the Fisher Information $I(\theta)$, i.e.,

$$\mathrm{Var}(\hat{\theta}) \geq \frac{1}{\mathrm{I}(\theta)}.$$

This bound is structure-aware and dependent on $\widetilde{\boldsymbol{\pi}}$ and the comparison graph structure, and hence not directly computable in closed-form here. In the paper, they conduct an empirical comparison between the Cramer-Rao lower bound and the error rate achieved by Rank Centrality and MLE (See Figure 3), and show that they match.

References

Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th International Conference on World Wide Web*, WWW '01, page 613–622, New York, NY, USA, 2001. Association for Computing Machinery. ISBN 1581133480. doi: 10.1145/371920.372165. URL https://doi.org/10.1145/371920.372165.

M. G. Kendall. Further contributions to the theory of paired comparisons. *Biometrics*, 11(1):43–62, 1955. ISSN 0006341X, 15410420. URL http://www.jstor.org/stable/3001479.

John G. Seeley. The net of reciprocal influence; a problem in treating sociometric data. *Canadian Journal of Psychology/Revue Canadienne de Psychologie*, 3(4):234–240, may 1949. doi: 10.1037/h0084096. URL https://doi.org/10.1037/h0084096.

A Multinomial Logit and Plackett-Luce Model

Multinomial Logit (MNL) Model. Multinomial Logit is a model for random choice behavior among more than two options. Concretely, given n items with weights $\{w_i\}$, an item i is assumed to be selected with probability

$$\Pr[\text{Choose item } i] = \frac{w_i}{\sum_{j\in[n]} w_j},$$

which is the same as sampling from a multinomial distribution with parameters $(w_1, \ldots, w_n)$. The "logit" part is due to the common reparameterization of $\theta_i = \log w_i$ to relax the assumption of $w_i > 0$ on the weights to $\theta_i \in [-\infty, \infty]$.

The MNL model can be seen to induce a distribution over full rankings; however, in ranking contexts, we are interested in sampling a full ordering that is faithful to the original distribution rather than only one item. The Plackett-Luce construction is an approach to sample an entire permutation while preserving MNL-like probabilities for each item.

Plackett-Luce Model. The Plackett-Luce model offers a way to generate a randomized permutation of the n itmes according to their weights $\{w_i\}$. See Algorithm 1 for pseudocode.

Algorithm 1 Sampling Full Ordering via Plackett-Luce Model

Require: $\{w_1, w_2, \ldots, w_n\}, \forall i \in [n] : w_i > 0$.

$\mathcal{U} \leftarrow \{1, 2, \ldots, n\}, \pi \leftarrow [\,]$

for $k = 1$ to n **do**

  $i(k) \sim \text{Categorical}\Big(\frac{w_1}{\sum_{j\in\mathcal{U}} w_j}, \frac{w_2}{\sum_{j\in\mathcal{U}} w_j}, \ldots, \frac{w_{|\mathcal{U}|}}{\sum_{j\in\mathcal{U}} w_j}\Big)$

  $\pi[k] \leftarrow i(k)$

  $\mathcal{U} \leftarrow \mathcal{U} \setminus i(k)$

end for

return $\pi = \big(i(1), i(2), \ldots, i(n)\big)$

Notice that picking the k-th item with the corresponding probabilities is consistent with the MNL principle for the remaining items at each step, and is also referred to as a *multistage MNL*.

It turns out that solving for the weights from a sampled set of permutations in an MLE manner cannot be done in closed form and requires iterative optimization. Namely, suppose we observe T independent permutations $\{\pi(1), \pi(2), \ldots, \pi(T)\}$, each drawn from the same Plackett-Luce distribution with unknown parameters $\{w_i\}$. We write the permutation for each sample t as $\pi(t) = \big(\pi_1(t), \pi_2(t), \ldots, \pi_n(t)\big)$. Then, the probability of observing this

ordering under the Plackett-Luce model is given by

$$\Pr\big[\pi(t)\big] = \prod_{k=1}^{n} \frac{w_{\pi_k(t)}}{\sum_{j=k}^{n} w_{\pi_j(t)}}.$$

As each permutation factorizes in this way, the likelihood of observing all T permutations is

$$L\Big(\{w_i\} \,\Big|\, \pi(1), \pi(2), \ldots, \pi(T)\Big) = \prod_{t=1}^{T} \prod_{k=1}^{n} \frac{w_{\pi_k(t)}}{\sum_{j=k}^{n} w_{\pi_j(t)}}.$$

Working with the log-likelihood, we obtain:

$$\log L = \sum_{t=1}^{T} \sum_{k=1}^{n} \left[\log w_{\pi_k(t)} - \log \left(\sum_{j=k}^{n} w_{\pi_j(t)} \right) \right].$$

The parameters $\{w_i\}$ appear both inside and outside the logarithms, and no closed-form solution is available.

Connection to BTL Model. While the Plackett-Luce model considers a comparison (ranking) over all $n-k$ items at the k-th time step, the BTL model instead considers pairwise rankings. This can be thought of as instead considering the probability that item j is ranked higher than i marginalizing over all other items, which is exactly equal to the expressions obtained in the BTL model. This is important, as if we assume the underlying discrete choice process is follows the Plackett–Luce model, then focusing only on pairwise outcomes is sufficient to glean the same parameters $\{w_i\}$.

B Other Models for Ranking

B.1 Mallows Model

The Mallows model is a distance-based probabilistic model over full rankings that assumes the existence of a true ranking and that other observed rankings are more or less likely depending on their dissimilarity to the central ranking.

Concretely, the probability of an observed ranking (permutation) σ is given by

$$\Pr(\sigma) \propto \exp\Big(-\theta \cdot d\big(\sigma, \sigma^*\big)\Big),$$

where σ^* is the central ranking, $\theta > 0$ is a dispersion parameter where larger θ concentrates probability nearer to σ^*, and $d(\cdot,\cdot)$ is a distance metric between rankings, typically the Kendall-τ distance (number of pairs $\{i, j\}$ such that the order of i and j in σ is different from σ^*).

In contrast to the BTL model, the Mallows model assigns likelihoods to complete rankings rather than to individual pairwise outcomes. Although one can derive pairwise marginals from the Plackett-Luce model, the Mallows model cannot yield pairwise comparison probabilities. Unfortunately, it turns out that MLE under the Mallows model is provably difficult (NP-hard).

Kemeny Optimization. One approach under the Mallows model is *Kemeny optimization*: Given a collection of (possibly inconsistent) rankings (or pairwise comparisons), the Kemeny optimal ranking is the one that minimizes the total disagreement with the inputs. In particular, Kemeny optimization seeks the ranking σ that minimizes

$$\sum_{\text{comparisons}} d\big(\sigma, \sigma(\text{comparisons})\big),$$

where $d(\cdot, \cdot)$ is the Kendall-τ distance. This formulation is equivalent to finding the ranking that best "matches" the pairwise preferences in a least-squares/minimum total inversion count sense. However, Kemeny optimization is known to be NP-hard.

In general, practitioners generally avoid the Mallows framework due to its computational intractability, whereas BTL is easier to work with.

C Analysis for Borda Count

C.1 Bad Case Construction

Assume n is even for simplicity. We consider a partitioning of the items into two equal-sized groups $S, T \subset [n]$. Consider true weights according to the BTL model as:

$$w_i = \begin{cases} 1, & \text{if } i \in S, \\ \epsilon, & \text{if } i \in T, \end{cases}$$

for some $1 \gg \epsilon > 0$. Then, for items i, j where $i \in S, j \in T$, the probability that i wins is

$$\Pr[S \succ T] = \frac{1}{1+\epsilon} \approx 1,$$

i.e., every item in S is much stronger than any item in T. Now suppose an adversary is allowed to choose which pairs to compare. Consider the following sampling design:

- Every pair (i, j) with $i, j \in S$ is compared k times
- Every pair (i, j) with $i, j \in T$ is compared k times
- No comparisons are made between any pairs (i, j) where $i \in S, j \in T$.

Here, the comparison graph is disconnected into two disjoint cliques S and T. Under the BTL model, the normalized true scores are given by:

$$\tilde{w}_i = \frac{w_i}{\sum_{j=1}^n w_j} = \begin{cases} \dfrac{1}{\frac{n}{2}(1+\epsilon)} = \dfrac{2}{n(1+\epsilon)} \approx \dfrac{2}{n}, & \text{if } i \in S, \\ \dfrac{\epsilon}{\frac{n}{2}(1+\epsilon)} = \dfrac{2\epsilon}{n(1+\epsilon)} \approx \dfrac{2\epsilon}{n}, & \text{if } i \in T. \end{cases}$$

Note that for $\epsilon \ll 1$, the strong items have a significantly larger normalized score than the weak items. Then, Borda count recovers the weights as follows:

- For any $i \in S$: All comparisons are within S, and since all items in S have the same true weight, the win probability for any pair within S is
$$\frac{1}{1+1} = \frac{1}{2}.$$
With many repeated comparisons, the empirical win fraction converges to $\frac{1}{2}$.
- Similarly, for any $j \in T$, Comparisons within T yield win probabilities of $\frac{1}{2}$.

Thus, in the limit of many comparisons (i.e. $k \to \infty$), the Borda count scores become

$$\text{Borda Score}(i) = \frac{1}{2}$$

for all $i \in S \cup T$ to obtain $\pi_{\text{Borda Count}}(i) = \frac{1}{n}$ for all $i \in [n]$. If we compute the relative error between the estimator π_B and the true vector $\tilde{w}$, it turns out that this becomes

$$\frac{\|\pi_B - \tilde{w}\|}{\|\tilde{w}\|} \approx \sqrt{\frac{1/n}{2/n}} = \sqrt{\frac{1}{2}},$$

which is a constant independent of k (and n).