
Fast and Accurate Ranking Regression

İlkay Yıldız
Dept. of ECE
Northeastern Univ.
yildizi@ece.neu.edu

Jennifer Dy
Dept. of ECE
Northeastern Univ.
jdy@ece.neu.edu

Deniz Erdoğan
Dept. of ECE
Northeastern Univ.
erdogmus@ece.neu.edu

Jayashree Kalpathy-Cramer
Dept. of Radiology
MGH/Harvard Medical School
kalpathy@nmr.mgh.harvard.edu

Susan Ostmo
Dept. of Ophthalmology
Casey Eye Inst., OHSU
ostmo@ohsu.edu

J. Peter Campbell
Dept. of Ophthalmology
Casey Eye Inst., OHSU
campbelp@ohsu.edu

Michael F. Chiang
Dept. of Ophthalmology
Casey Eye Inst., OHSU
chiangm@ohsu.edu

Stratis Ioannidis
Dept. of ECE
Northeastern Univ.
ioannidis@ece.neu.edu

Abstract

We consider a ranking regression problem in which we use a dataset of ranked choices to learn Plackett-Luce scores as functions of sample features. We solve the maximum likelihood estimation problem by using the Alternating Directions Method of Multipliers (ADMM), effectively separating the learning of scores and model parameters. This separation allows us to express scores as the stationary distribution of a continuous-time Markov Chain. Using this equivalence, we propose two spectral algorithms for ranking regression that learn model parameters up to 579 times faster than the Newton’s method.

1 Introduction

Learning from ranked choices has a long history in domains such as econometrics (McFadden, 1973; Ryzin and Mahajan, 1999), transportation (McFadden, 2000), psychometrics (Thurstone, 1927; Bradley and Terry, 1952), and sports (Elo, 1978), to name a few. The Plackett-Luce choice model (Plackett, 1975) is a popular parametric model used for inference in this setting: each sample is parametrized by a score and the probability that a sample is ranked higher than a set of alternatives is proportional to this score.

Plackett-Luce scores are traditionally learned from ranking observations via Maximum Likelihood Estimation (MLE) (Dykstra, 1960; Hunter, 2004; Hajek et al., 2014; Negahban et al., 2018); under a reparametrization, the negative log-likelihood is convex

and Plackett-Luce scores can be estimated via, e.g., Newton’s method (Nocedal and Wright, 2006). Nevertheless, for large datasets, Newton’s method can be prohibitively slow (Hunter, 2004). Recently, Maystre and Grossglauser (2015) proposed a highly efficient iterative spectral method, termed Iterative Luce Spectral Ranking (ILSR), that estimates Plackett-Luce scores significantly faster than state-of-the-art methods. ILSR relies on the fact that ML estimates of Plackett-Luce scores constitute the stationary distribution of a Markov chain with transition rates defined by ranking observations.

The above approaches learn Plackett-Luce scores in the absence of sample features, which precludes rank predictions on samples outside the training set. A natural variant of the above setting is *ranking regression*, whereby Plackett-Luce scores are parametrized functions of sample features. This problem has received considerable attention in the literature, via both shallow (Joachims, 2002; Pahikkala et al., 2009; Tian et al., 2019) and deep models (Burgess et al., 2005; Chang et al., 2016; Dubey et al., 2016; Han, 2018; Yıldız et al., 2019). Nevertheless, virtually all existing work on ranking regression relies on classic optimization methods for parameter inference. To the best of our knowledge, the opportunity to accelerate learning in ranking regression via spectral methods has not yet been explored.

We make the following contributions.

- We solve the ranking regression problem by using the Alternating Directions Method of Multipliers (ADMM) (Boyd et al., 2011) to perform MLE, effectively separating the learning of scores and model parameters. This separation allows us to express scores as the stationary distribution of a modified Markov Chain, and to devise spectral algorithms for ranking regression akin to ILSR.
- In particular, we propose two iterative algorithms, PLADMM and PLADMM-log, that jointly estimate

model parameters and Plackett-Luce scores via a spectral method. Though the problems solved are non-convex, we establish conditions that yield convergence guarantees, as well as initializations tailored to the Plackett-Luce objective.

- Our algorithms yield significant performance dividends in terms of both speed and accuracy on synthetic and real-life datasets. PLADMM and PLADMM-log are up to 579 times faster than traditional optimization methods regressing Plackett-Luce scores from features, including Newton’s method. Furthermore, for large datasets, PLADMM and PLADMM-log outperform feature-less methods, including ILSR, by 13% in maximal choice prediction accuracy and 9% in ranking prediction Kendall-Tau correlation.

From a technical standpoint, we show that the Plackett-Luce negative log-likelihood *augmented with a proximal penalty* has stationary points that satisfy the balance equations of a Markov chain (c.f. Thm 4.2). In turn, ADMM allows us to reduce ranking regression to a regularized MLE with precisely such a penalty. The remainder of this paper is organized as follows. We review related literature in Sec. 2. We formulate our problem in Sec. 3 and summarize ILSR. We describe our main contributions and proposed algorithms in Sec. 4. We present our experiments in Sec. 5 and conclude with future work in Sec. 6.

2 Related Work

The problem of *rank aggregation* (Dwork et al., 2001), in which a total ordering of samples is regressed from ranking observations, is classic; literature on the subject is vast—see, e.g., the surveys by Fligner and Verducci (1993), Cattelan (2012) and Marden (2014). Probabilistic inference in this setting typically assumes (a) that a “true” total ordering of samples exists, and (b) that ranking observations exhibit a *stochastic transitivity property* (Agarwal, 2016): a sample is more likely to be ranked higher than another when this event is consistent with the underlying total ordering.

The noisy permutation model is a non-parametric model for this setting: pairwise comparisons consistent with the underlying total ordering are observed under i.i.d. Bernoulli noise. Maximum likelihood estimation (MLE) is NP-hard in this setting. A polytime algorithm by Braverman and Mossel (2008) recovers the underlying ordering in $\Theta(n \log n)$ comparisons, w.h.p.; this is tightened by several recent works (Wauthier et al., 2013; Mao et al., 2017, 2018). The Mallows model (Mallows, 1957) assumes that the probability of a ranking observation is a decreasing function of its distance from the underlying total ordering, under appropriate notions of distance (e.g., Kendall-Tau); MLE can be approached, e.g., via EM (Lu and Boutilier, 2011).

Shah et al. (2016a) learn the full matrix of pairwise comparison probabilities via a minimax optimal estimator requiring $\Theta(\log^2 n)$ comparisons. Rajkumar and Agarwal (2016) learn the matrix via matrix completion, requiring $\Theta(nr \log n)$ comparisons, where $r \ll n$ is the rank. Ammar and Shah (2011) assume that comparisons are sampled from an unknown distribution over total orderings and propose an entropy maximization algorithm requiring $\Theta(n^2)$ comparisons.

We focus on parametric models, as they are more natural in the context of regressing rankings from sample features. In both Plackett-Luce (Plackett, 1975) and Thurstone (Thurstone, 1927) each sample is parametrized by a score. In the Thurstone model, observations result from comparing scores after the addition of Gaussian noise. Vojnovic and Yun (2016) and Shah et al. (2016b) estimate Thurstone scores via MLE and provide sample complexity bounds that are inversely proportional to the smallest non-zero eigenvalue of the Laplacian of a graph modeling comparisons. In Plackett-Luce, the probability that a sample is chosen over a set of alternatives is proportional to its score. Hunter (2004) proposes a Minorization-Maximization (MM) approach to estimate Plackett-Luce scores via MLE, earlier used by Dykstra (1960) on pairwise comparisons (i.e., on the Bradley-Terry (BT) setting). Hajek et al. (2014) provide an upper bound on the error in estimating the Plackett-Luce scores via MLE and show that the latter is minimax-optimal. Negahban et al. (2018) propose a latent factor model, estimating parameters via a convex relaxation of the corresponding rank penalty and providing sample complexity guarantees. Assuming score priors, Guiver and Snelson (2009), Caron and Doucet (2012) and Azari et al. (2012) estimate Plackett-Luce scores via Bayesian inference.

Our focus on Plackett-Luce is due to the recent emergence of spectral algorithms for inference in this setting. Negahban et al. (2012) propose the Rank Centrality (RC) algorithm for the BT setting and derive a minimax error bound. Chen and Suh (2015) propose a spectral MLE algorithm extending RC with an additional stage that cyclically performs MLE for each score. Soufiani et al. (2013) and Jang et al. (2017) extend RC to rankings of two or more samples by breaking rankings into independent comparisons. Improved bounds, applying also to broader noise settings, are provided by Rajkumar and Agarwal (2014). Khetan and Oh (2016) generalize the work by Soufiani et al. (2013) by breaking rankings into independent shorter rankings, and building a hierarchy of tractable and consistent estimators. Blanchet et al. (2016) model sequential choices by state transitions in a Markov chain (MC), where transitions are functions of choice probabilities.

Bridging the above approaches with MLE, Maystre and

Grossglauser (2015) show that the MLE of Plackett-Luce scores can be expressed as the stationary distribution of an MC. Their proposed Iterative Luce Spectral Ranking (ILSR) algorithm estimates the Plackett-Luce scores faster than traditional optimization methods, such as, e.g., Hunter’s (Hunter, 2004) and Newton’s method, and more accurately than prior spectral rank aggregation methods. Ragain and Ugander (2016) show that a spectral approach applies even after relaxing the assumption that the relative order of any two samples is independent of the alternatives. Agarwal et al. (2018) propose another spectral method called accelerated spectral ranking by departing from the exact equivalence between MLE and MC approximation and demonstrate faster convergence than ILSR.

We depart from all aforementioned methods by regressing ranked choices from sample features. Closer to our work, RankSVM (Joachims, 2002) learns a target ranking from features via a linear Support Vector Machine (SVM), with constraints imposed by all possible comparisons. Pahikkala et al. (2009) propose a regularized least-squares based algorithm for learning to rank from comparisons. Several works learn comparisons from features via MLE over logistic BT models (Guo et al., 2018; Tian et al., 2019); deeper models have also been considered (Burgess et al., 2005; Chang et al., 2016; Dubey et al., 2016; Han, 2018; Yıldız et al., 2019). Niranjana and Rajkumar (2017) assume that features are low-dimensional and use matrix completion to recover the BT scores. Saha and Rajkumar (2018) propose a least squares based algorithm called f-BTL to regress the BT scores; we adjust this to initialize our algorithm. To the best of our knowledge, we are the first to use a spectral method akin to ILSR to (a) regress Plackett-Luce scores from features, and (b) to establish a significant speedup over prior art.

3 Problem Formulation

Plackett-Luce Model. We consider a dataset of n samples indexed by $i \in \mathcal{N} \equiv \{1, \dots, n\}$. Every sample $i \in \mathcal{N}$ has a corresponding p -dimensional feature vector $\mathbf{x}_i \in \mathbb{R}^p$. There exists an underlying total ordering of these n samples. A labeler of this dataset acts as a (possibly noisy) oracle revealing this total ordering: when presented with a query $A \subseteq \mathcal{N}$, i.e., a set of alternative samples, the noisy labeler chooses the *maximal* sample in A w.r.t. the underlying total ordering.

Formally, our “labeled” dataset $\mathcal{D} = \{(c_\ell, A_\ell) \mid \ell \in \mathcal{M} = \{1, \dots, M\}\}$ consists of M observations (c_ℓ, A_ℓ) , $\ell \in \mathcal{M}$, where $A_\ell \subseteq \mathcal{N}$ is the ℓ -th query submitted to the labeler and $c_\ell \in A_\ell$ is her respective ℓ -th maximal choice (i.e., the label). We tackle the problem of regressing such choices c_ℓ from the features \mathbf{x}_i of the samples $i \in A_\ell$. To do so, we assume that choices are governed by

the Plackett-Luce model (Plackett, 1975). The model asserts that every sample $i \in \mathcal{N}$ is associated with a non-negative deterministic score $\pi_i \in \mathbb{R}_+$. Given scores $\boldsymbol{\pi} = [\pi_i]_{i \in \mathcal{N}} \in \mathbb{R}_+^n$, then (a) observations (c_ℓ, A_ℓ) , $\ell \in \mathcal{M}$ are independent, and (b) given query A_ℓ ,

$$\mathbf{P}(c_\ell \mid A_\ell, \boldsymbol{\pi}) = \pi_{c_\ell} / \sum_{j \in A_\ell} \pi_j = \pi_{c_\ell} / \sum_{j \in A_\ell} \pi_j. \quad (1)$$

Abusing notation, we write the score of the chosen sample as $\pi_\ell \equiv \pi_{c_\ell}$. Note that $\mathbf{P}(c_\ell \mid A_\ell, \boldsymbol{\pi}) = \mathbf{P}(c_\ell \mid A_\ell, s\boldsymbol{\pi})$, for all $s > 0$; thus, w.l.o.g., we may additionally assume (or enforce via rescaling) that Plackett-Luce scores satisfy $\mathbf{1}^\top \boldsymbol{\pi} = 1$.

Plackett-Luce also applies to *ranking* data. In the ranking setting, when presented with a query $A_\ell \subseteq \mathcal{N}$, the labeler ranks the samples in A_ℓ into an ordered sequence $\alpha_1^\ell > \alpha_2^\ell > \dots > \alpha_{|A_\ell|}^\ell$. Under the Plackett-Luce model, this ranking is expressed as $|A_\ell| - 1$ maximal choice queries: α_1^ℓ over A_ℓ , α_2^ℓ over $A_\ell \setminus \{\alpha_1^\ell\}$, etc., so that:

$$\mathbf{P}(\alpha_1^\ell > \alpha_2^\ell > \dots > \alpha_{|A_\ell|}^\ell \mid A_\ell, \boldsymbol{\pi}) = \prod_{t=1}^{|A_\ell|-1} (\pi_{\alpha_t^\ell} / \sum_{s=t}^{|A_\ell|} \pi_{\alpha_s^\ell}). \quad (2)$$

The product form of (2) implies that rankings of a query A_ℓ can be converted to $|A_\ell| - 1$ maximal-choice observations, each governed by (1), that have the same joint probability: ranking $(\alpha_1^\ell > \alpha_2^\ell > \dots > \alpha_{|A_\ell|}^\ell)$ can be seen as the outcome of α_1^ℓ being chosen as the top within the query set A_ℓ , α_2^ℓ being the top among $A_\ell \setminus \{\alpha_1^\ell\}$, etc. Keeping this reduction from ranking to maximal-choice datasets in mind, we focus on the latter in our exposition below.

Parameter Inference and Regression. Given observations \mathcal{D} , Maximum Likelihood Estimation (MLE) of the Plackett-Luce scores $\boldsymbol{\pi} \in \mathbb{R}_+^n$ amounts to minimizing the negative log-likelihood:

$$\mathcal{L}(\mathcal{D} \mid \boldsymbol{\pi}) \equiv \sum_{\ell=1}^M (\log \sum_{j \in A_\ell} \pi_j - \log \pi_\ell). \quad (3)$$

To regress scores $\boldsymbol{\pi}$ from sample features \mathbf{x}_i , $i \in \mathcal{N}$, we consider two cases:

Affine Case. We assume that there exist $\boldsymbol{\beta} \in \mathbb{R}^p$ and $b \in \mathbb{R}$ such that $\boldsymbol{\pi} = \boldsymbol{\pi}_{\text{AFF}}(\boldsymbol{\beta}, b; \mathbf{X}) \equiv \mathbf{X}\boldsymbol{\beta} + b\mathbf{1}$. Then, MLE of parameters $(\boldsymbol{\beta}, b) \in \mathbb{R}^{p+1}$ amounts to solving:

$$\min_{(\boldsymbol{\beta}, b): \boldsymbol{\pi}_{\text{AFF}}(\boldsymbol{\beta}, b; \mathbf{X}) \geq \mathbf{0}} \mathcal{L}(\mathcal{D} \mid \boldsymbol{\pi}_{\text{AFF}}(\boldsymbol{\beta}, b; \mathbf{X})), \quad (4)$$

where \mathcal{L} is given by (3), and $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T \in \mathbb{R}^{n \times p}$. We note that Problem (4) is *not* convex, as the objective is not convex in $(\boldsymbol{\beta}, b) \in \mathbb{R}^{p+1}$ (c.f. Appendix A).

Logistic Case. In the logistic case, we assume that there exists $\boldsymbol{\beta} \in \mathbb{R}^p$ s.t. $\boldsymbol{\pi} = \boldsymbol{\pi}_{\text{LOG}}(\boldsymbol{\beta}; \mathbf{X}) \equiv [e^{\boldsymbol{\beta}^\top \mathbf{x}_i}]_{i \in \mathcal{N}}$. As $\boldsymbol{\pi}_{\text{LOG}}(\boldsymbol{\beta}; \mathbf{X}) \geq \mathbf{0}$ by definition, MLE corresponds to:

$$\min_{\boldsymbol{\beta} \in \mathbb{R}^p} \mathcal{L}(\mathcal{D} \mid \boldsymbol{\pi}_{\text{LOG}}(\boldsymbol{\beta}; \mathbf{X})). \quad (5)$$

The objective of (5) is convex in β (c.f. Appendix A), so an optimal solution can be found via, e.g., Newton’s method (Nocedal and Wright, 2006).

Plackett-Luce Without Features and a Spectral Method. We wish to construct highly efficient algorithms for solving regression problems (4) and (5). To do so, we first briefly review the state of the art for learning the scores π in the absence of features. In this case, MLE amounts to:

$$\min_{\pi \in \mathbb{R}_+^n} \mathcal{L}(\mathcal{D} | \pi). \quad (6)$$

As is the case for (5), reparametrizing the scores as $\pi_i = e^{\theta_i}$, $i \in \mathcal{N}$ makes the negative log-likelihood \mathcal{L} convex in $\theta = [\theta_i]_{i \in \mathcal{N}}$, which in turn enables computing the Plackett-Luce scores via Newton’s method. Nevertheless, Newton’s method can be prohibitively slow for large n and M (Hunter, 2004). Recently, Maystre and Grossglauser (2015) proposed a novel spectral algorithm that is significantly faster than Newton’s method. Their algorithm relies on the following theorem which, for completeness, we re-prove in Appendix B:

Theorem 3.1 (Maystre and Grossglauser (2015)). *An optimal solution $\pi \in \mathbb{R}_+^n$ to (6) satisfies:*

$$\sum_{j \neq i} \pi_j \lambda_{ji}(\pi) = \sum_{j \neq i} \pi_i \lambda_{ij}(\pi), \quad \text{for all } i \in \mathcal{N}, \quad (7)$$

where, for all $i, j \in \mathcal{N}$, with $i \neq j$,

$$\lambda_{ji}(\pi) = \sum_{\ell \in W_i \cap L_j} \left(\sum_{t \in A_\ell} \pi_t \right)^{-1} \geq 0, \quad (8)$$

for $W_i = \{\ell | i \in A_\ell, c_\ell = i\}$ the observations where sample $i \in \mathcal{N}$ is chosen and $L_i = \{\ell | i \in A_\ell, c_\ell \neq i\}$ the observations where sample $i \in \mathcal{N}$ is not chosen.

Eq. (7) are the *balance equations* of a continuous-time Markov Chain (MC) with transition rates:

$$\mathbf{\Lambda}(\pi) = [\lambda_{ji}(\pi)]_{i,j \in \mathcal{N}}, \quad (9)$$

where $\lambda_{ji}(\pi)$ are given by Eq. (8). Hence, π is the stationary distribution of the MC defined by transition rates $\mathbf{\Lambda}(\pi)$ (Gallager, 2013). Let $\text{ssd}(\mathbf{\Lambda})$ be the stationary distribution of an MC with transition rates $\mathbf{\Lambda}$. When matrix $\mathbf{\Lambda}$ is fixed (i.e., the transition rates are known), the vector $\text{ssd}(\mathbf{\Lambda})$ is a solution to the linear system defined by the balance equations (7) and $\mathbf{1}^\top \pi = 1$, as it is a distribution.¹ If (9) is irreducible, the linear system has a unique solution $\pi > \mathbf{0}$ (Gallager, 2013).

However, the transition matrix $\mathbf{\Lambda} = \mathbf{\Lambda}(\pi)$ in Theorem 3.1 is itself a function of π , and is therefore a priori unknown. Maystre and Grossglauser (2015) find

¹In practice, $\text{ssd}(\mathbf{\Lambda})$ can be computed by uniformizing $\mathbf{\Lambda}$, i.e., increasing self-transition rates until all states have the same outgoing rate, and finding the leading left eigenvector via, e.g., the power method (Lei et al., 2016).

π through an iterative algorithm. Starting from the uniform distribution $\pi^0 = \frac{1}{n} \mathbf{1}$, they compute:

$$\pi^{l+1} = \text{ssd}(\mathbf{\Lambda}(\pi^l)), \quad \text{for } l = 0, 1, 2, \dots, \quad (10)$$

where $\mathbf{\Lambda}(\cdot)$ is given by (8), (9). Maystre and Grossglauser (2015) refer to Eq. (10) as the *Iterative Luce Spectral Ranking* (ILSR) algorithm. They also establish that (10) converges to an optimal solution of (6) under mild assumptions. Most importantly, as mentioned above, ILSR significantly outperforms state-of-the-art MLE algorithms in computational efficiency.

4 Plackett-Luce ADMM (PLADMM) Algorithm

Given ILSR’s significant computational benefits, we wish to develop analogues in the regression setting. In contrast to the feature-less setting, it is not a priori evident how to solve Problems (4) and (5) via a spectral approach. Taking the affine case as an example, and momentarily ignoring issues of non-convexity, the stationary points of the Lagrangian of the optimization problem (4) cannot be expressed via the balance equations of an MC. Our main contribution is to circumvent this problem by using the Alternating Directions Method of Multipliers (ADMM) (Boyd et al., 2011). Intuitively, ADMM allows us to decouple the optimization of scores π from model parameters β and b , encapsulating them in a quadratic penalty: the latter becomes amenable to a spectral approach after a series of manipulations that we outline below (see Thm. 4.2). We focus here on the affine case, extending our method to the logistic case in Appendix E.

An ADMM Approach. We rewrite Problem (4) as:

$$\text{Minimize } \mathcal{L}(\mathcal{D} | \pi) \quad (11a)$$

$$\text{subject to: } \pi = \mathbf{X}\beta + b\mathbf{1}, \quad \pi \geq \mathbf{0}. \quad (11b)$$

To simplify our notation, we introduce $\tilde{\beta} = (\beta, b) \in \mathbb{R}^{p+1}$ and $\tilde{\mathbf{X}} = [\mathbf{X} | \mathbf{1}] \in \mathbb{R}^{n \times (p+1)}$, so that $\pi = \tilde{\mathbf{X}}\tilde{\beta}$. ADMM solves (11) by minimizing the following *augmented Lagrangian*:

$$\begin{aligned} L_\rho(\tilde{\beta}, \pi, \mathbf{y}) &= \mathcal{L}(\mathcal{D} | \pi) + \mathbf{y}^\top (\tilde{\mathbf{X}}\tilde{\beta} - \pi) \\ &\quad + \frac{\rho}{2} \|\tilde{\mathbf{X}}\tilde{\beta} - \pi\|_2^2, \end{aligned} \quad (12)$$

where $\mathbf{y} \in \mathbb{R}^n$ is a dual variable corresponding to the equality constraints in Eq. (11) and $\rho > 0$ is a penalty parameter. ADMM alternates between optimizing $\tilde{\beta}$ and π , thereby decoupling these two variables. Using a rescaling $\mathbf{u} = \frac{1}{\rho} \mathbf{y} \in \mathbb{R}^n$ for convenience, applying ADMM on problem (11) yields the following iterative algorithm (see Appendix C for a detailed derivation):

$$\tilde{\beta}^{k+1} = \arg \min_{\tilde{\beta} \in \mathbb{R}^{p+1}} \|\tilde{\mathbf{X}}\tilde{\beta} - \pi^k + \mathbf{u}^k\|_2^2, \quad (13a)$$

$$\boldsymbol{\pi}^{k+1} = \arg \min_{\boldsymbol{\pi} \in \mathbb{R}_+^n} (\mathcal{L}(\mathcal{D}|\boldsymbol{\pi}) + \frac{\rho}{2} \|\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}^{k+1} - \boldsymbol{\pi} + \mathbf{u}^k\|_2^2), \quad (13b)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}^{k+1} - \boldsymbol{\pi}^{k+1}. \quad (13c)$$

This has the following immediate computational advantages. First, step (13a) is a quadratic minimization and admits a closed form solution. Crucially, step (13b) is amenable to a spectral approach, though the corresponding MC is not as apparent as in ILSR; we outline its construction below.

An MC for Step (13b). We first establish the following auxiliary lemma, proved in Appendix D.1.

Lemma 4.1. *Given $\tilde{\boldsymbol{\beta}}^{k+1} \in \mathbb{R}^{p+1}$, $\mathbf{u}^k \in \mathbb{R}^n$, let $\boldsymbol{\pi} \in \mathbb{R}_+^n$ be such that:*

$$\nabla_{\boldsymbol{\pi}} L_{\rho}(\tilde{\boldsymbol{\beta}}^{k+1}, \boldsymbol{\pi}, \mathbf{u}^k) = \mathbf{0}. \quad (14)$$

For $\boldsymbol{\sigma} = [\sigma_i]_{i \in \mathcal{N}} \equiv \rho(\boldsymbol{\pi} - \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}^{k+1} - \mathbf{u}^k)$, and $[\lambda_{ij}(\boldsymbol{\pi})]_{i,j \in \mathcal{N}}$ given by (8), (14) is equivalent to:

$$\sum_{j \neq i} \pi_j \lambda_{ji}(\boldsymbol{\pi}) - \sum_{j \neq i} \pi_i \lambda_{ij}(\boldsymbol{\pi}) = \pi_i \sigma_i, \quad (15)$$

for all $i \in \mathcal{N}$.

Although Eq. (15) looks similar to Eq. (7), it is not evident that it corresponds to the balance equations of an MC as, in general, $\boldsymbol{\sigma} \neq \mathbf{0}$. Nevertheless, we prove that this is indeed the case:

Theorem 4.2. *Eq. (15) are the balance equations of a continuous-time MC with transition rates:*

$$\mu_{ji}(\boldsymbol{\pi}) = \begin{cases} \lambda_{ji}(\boldsymbol{\pi}) + \frac{2\pi_i \sigma_i \sigma_j}{\sum_{t \in \mathcal{N}_-} \pi_t \sigma_t - \sum_{t \in \mathcal{N}_+} \pi_t \sigma_t} & \text{if } j \in \mathcal{N}_+ \text{ and } i \in \mathcal{N}_- \\ \lambda_{ji}(\boldsymbol{\pi}) & \text{otherwise,} \end{cases} \quad (16)$$

where $\boldsymbol{\sigma} = [\sigma_i]_{i \in \mathcal{N}} \equiv \rho(\boldsymbol{\pi} - \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}^{k+1} - \mathbf{u}^k)$, $[\lambda_{ij}(\boldsymbol{\pi})]_{i,j \in \mathcal{N}}$ are given by (8), and $(\mathcal{N}_+, \mathcal{N}_-)$ is a partition of \mathcal{N} such that $\sigma_i \geq 0$ for all $i \in \mathcal{N}_+$ and $\sigma_i < 0$ for all $i \in \mathcal{N}_-$.

The proof is in App. D.2. By Lemma 4.1 and Theorem 4.2, we conclude that a stationary $\boldsymbol{\pi} \in \mathbb{R}_+^n$ satisfying (14) is also the stationary distribution of the continuous-time MC with transition rates:

$$\mathbf{M}(\boldsymbol{\pi}) = [\mu_{ji}(\boldsymbol{\pi})]_{i,j \in \mathcal{N}}, \quad (17)$$

where $\mu_{ji}(\boldsymbol{\pi})$ are given by Eq. (16). Motivated by these observations, and mirroring ILSR (Eq. (10)), we compute a solution to (13b) via:

$$\boldsymbol{\pi}^{l+1} = \text{ssd}(\mathbf{M}(\boldsymbol{\pi}^l)), \quad \text{for } l = 0, 1, 2, \dots, \quad (18)$$

where $\mathbf{M}(\cdot)$ is given by Eq. (17). We refer to this procedure as ILSRX (“ILSR with features”).

Overall Algorithm. Putting everything together, our Plackett-Luce ADMM (PLADMM) solving Eq. (11)

Algorithm 1 PLADMM

```

1: procedure ADMM( $\tilde{\mathbf{X}}$ ,  $\mathcal{D} = \{(c_\ell, A_\ell) \mid \ell \in \mathcal{M}\}$ ,  $\rho$ )
2:   Initialize  $\tilde{\boldsymbol{\beta}}$  via Eq. (20);  $\boldsymbol{\pi} \leftarrow \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}$ ;  $\mathbf{u} \leftarrow \mathbf{0}$ 
3:   repeat
4:      $\boldsymbol{\pi} \leftarrow \text{ILSRX}(\rho, \boldsymbol{\pi}, \tilde{\mathbf{X}}, \tilde{\boldsymbol{\beta}}, \mathbf{u})$ 
5:      $\mathbf{u} \leftarrow \mathbf{u} + \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} - \boldsymbol{\pi}$ 
6:      $\tilde{\boldsymbol{\beta}} \leftarrow (\tilde{\mathbf{X}}^T \tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}^T (\boldsymbol{\pi} - \mathbf{u})$ 
7:   until convergence
8: return  $\tilde{\boldsymbol{\beta}}, \boldsymbol{\pi}$ 
9: end procedure

1: procedure ILSRX( $\rho, \boldsymbol{\pi}, \tilde{\mathbf{X}}, \tilde{\boldsymbol{\beta}}, \mathbf{u}$ )
2:   repeat
3:      $\boldsymbol{\sigma} \leftarrow \rho(\boldsymbol{\pi} - \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} - \mathbf{u})$ 
4:     Calculate  $\mathbf{M}(\boldsymbol{\pi}) = [\mu_{ji}(\boldsymbol{\pi})]_{i,j \in \mathcal{N}}$  via Eq. (16)
5:      $\boldsymbol{\pi} \leftarrow \text{ssd}(\mathbf{M}(\boldsymbol{\pi}))$ 
6:   until convergence
7: return  $\boldsymbol{\pi}$ 
8: end procedure
    
```

is summarized in Algorithm 1. We iteratively update $\tilde{\boldsymbol{\beta}}$, $\boldsymbol{\pi}$, and \mathbf{u} via Eq. (13) until convergence, with $\tilde{\boldsymbol{\beta}}$ updated via Eq. (13a) and $\boldsymbol{\pi}$ updated via ILSRX (Eq. (18)). At iteration k , we initialize ILSRX with $\boldsymbol{\pi}^{k-1}$. We note that, as Problem (11) is non-convex, selecting a good initialization point is important in practice. We discuss initialization, additional computational issues, and theoretical guarantees below.

Initialization. We initialize $\tilde{\boldsymbol{\beta}}$ so that $\tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}}$ is a good approximation of Plackett-Luce scores. We use a technique akin to Saha and Rajkumar (2018), applied to our affine setting. Given a distribution over queries $A \subseteq \mathcal{N}$, let $P_{ij} = \mathbb{E}_A[c = i \mid \{i, j\} \subseteq A]$ be the probability that i is chosen given a query A that contains both i and j . By (1), for $i, j \in \mathcal{N}$, $\frac{P_{ij}}{P_{ji}} = \frac{\pi_i}{\pi_j} = \frac{\mathbf{x}_i^\top \boldsymbol{\beta} + b}{\mathbf{x}_j^\top \boldsymbol{\beta} + b}$, or:

$$\delta_{ij}(\tilde{\boldsymbol{\beta}}) \equiv (P_{ij} \mathbf{x}_j - P_{ji} \mathbf{x}_i)^\top \tilde{\boldsymbol{\beta}} + (P_{ij} - P_{ji})b = 0. \quad (19)$$

Motivated by (19), we estimate P_{ij} empirically from \mathcal{D} , and obtain our initialization $\tilde{\boldsymbol{\beta}}^0 = (\boldsymbol{\beta}^0, b^0) \in \mathbb{R}^{p+1}$ by solving (19) in the least-square sense; that is,

$$\tilde{\boldsymbol{\beta}}^0 = \arg \min_{\tilde{\boldsymbol{\beta}} \in \mathbb{R}^{p+1}: \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} \geq \mathbf{0} \wedge \mathbf{1}^\top \tilde{\mathbf{X}}\tilde{\boldsymbol{\beta}} = 1} \sum_{i,j} \delta_{ij}^2(\tilde{\boldsymbol{\beta}}). \quad (20)$$

Note that this is a convex quadratic program. Finally, we also set the initial dual variable as $\mathbf{u}^0 = \mathbf{0}$.

Computational Complexity. Each iteration of PLADMM involves the three steps in Eq. (13). One iteration of ILSRX is $O(\sum_{\ell \in \mathcal{D}} |A_\ell| + n^2)$ for constructing the transition matrix $\mathbf{M}(\boldsymbol{\pi})$ via Eq. (17) and for finding the stationary distribution $\boldsymbol{\pi}$ via, e.g., a power method (Lei et al., 2016), respectively. Updates of \mathbf{u} and $\tilde{\boldsymbol{\beta}}$ are both $O(n(p+1))$ as matrix-vector multiplications, since the matrix $(\tilde{\mathbf{X}}^T \tilde{\mathbf{X}})^{-1} \tilde{\mathbf{X}}^T$ can be precomputed.

Theoretical Guarantees. In general, condition (14) is not sufficient for optimality w.r.t. step (13b). To show this, we require the following technical assumption:

Assumption 4.1. *For $\{\boldsymbol{\pi}^k\}_{k \in \mathbb{N}}$, given by (13b), there exists an $\epsilon > 0$ such that $\pi_i^k > \epsilon$ for all $i \in \mathcal{N}$ and $k \in \mathbb{N}$.*

Spec.	Dataset			
	ROP	FAC	Pairwise Sushi	Triplet Sushi
n	100	1000	100	100
p	143	50	18	18
M	29705	728	450	1200
$ A_\ell $	2	2	2	3
n_{fold}	10	10	3	10
Type	Choice	Choice	Choice	Ranking

Table 1: No. of samples (n), no. of parameters (p), no. of observations (M), query size ($|A_\ell|$), no. of cross validation folds (n_{fold}), and type of observations for real data

Under this assumption, we show that stationarity implies optimality w.r.t. (13b) for large enough ρ :

Theorem 4.3. *Under Assumption 4.1, for $\rho \geq \frac{2}{\epsilon^2} \max_i \sum_{\ell \in A_\ell} \frac{1}{|A_\ell|^2}$, a $\boldsymbol{\pi} > \mathbf{0}$ satisfying condition (14) is a minimizer of (13b).*

The proof is in Appendix D.3. Moreover, although problem (11) is non-convex, we establish the following convergence guarantee for the ADMM steps (13). We provide the proof in Appendix D.4.

Theorem 4.4. *Suppose that there exists $\kappa > 0$ such that $\tilde{\mathbf{X}}^T \tilde{\mathbf{X}} \succeq \kappa \mathbf{I}$ and the sequence $\{(\boldsymbol{\pi}^k, \mathbf{u}^k, \tilde{\boldsymbol{\beta}}^{k+1})\}_{k \in \mathbb{N}}$ generated by (13) is bounded. Then, under Assumption 4.1, for $\rho > \frac{2 \max_i |W_i|}{\epsilon^2}$ where W_i is defined in Theorem 3.1, the sequence $\{(\boldsymbol{\pi}^k, \mathbf{u}^k, \tilde{\boldsymbol{\beta}}^{k+1})\}_{k \in \mathbb{N}}$ generated by (13) converges to a point that satisfies the Karush-Kuhn-Tucker (KKT) conditions of (11).*

5 Experiments

Experiment Setup. We evaluate PLADMM and PLADMM-log (the spectral algorithm for the logistic case, described in Appendix E) on synthetic and real-life datasets, summarized in Table 1. Additional details on our datasets are in Appendix F.1. We perform 10-fold cross validation (CV) for each dataset, except Pairwise Sushi, for which we use 3 folds. For synthetic datasets, we also repeat experiments over 5 random generations. We partition each dataset into training and test sets in two ways. In *observation CV*, we partition the dataset w.r.t. *observations* \mathcal{M} , using 90% of the M observations for training and the remaining 10% for testing. In *sample CV*, we partition samples \mathcal{N} , using 90% of the n samples for training and the remaining 10% for testing. When partitioning w.r.t. samples, *observations containing samples from both training and test partitions are discarded*. As the Pairwise Sushi dataset contains few observations (c.f. Table 1), we perform 3-fold cross validation in this case.

We implement² seven inference algorithms described in detail in Appendix F.2. Four are *feature methods*, i.e., algorithms that regress Plackett-Luce scores from features: PLADMM described in Algorithm 1,

PLADMM-log described in Appendix E, sequential least-squares quadratic programming (SLSQP), that solves (4), and Newton on $\boldsymbol{\beta}$, that solves the convex problem (5) via Newton’s method. The remaining three are *featureless methods*, i.e., algorithms that learn the Plackett-Luce scores from the choice observations alone: Iterative Luce Spectral Ranking (ILSR) described by Eq.(10), the Minorization-Maximization (MM) algorithm (Hunter, 2004), and Newton on $\boldsymbol{\theta}$ that solves Eq. (6) via the reparametrization $\pi_i = e^{\theta_i}$, $i \in \mathcal{N}$ using Newton’s method on $\boldsymbol{\theta} = [\theta_i]_{i \in \mathcal{N}}$.

Performance Metrics. We run each algorithm until convergence (see App. F.2 for criteria). We measure the elapsed time, including time spent in initialization, in seconds (Time) and the number of iterations (Iter). We measure the prediction performance by Top-1 accuracy (Top-1 Acc.) and Kendall-Tau correlation (KT) on the test set; formulas are provided in App. F.3. For synthetic datasets, we also measure the quality of convergence by the norm of the difference between estimated and true Plackett-Luce scores ($\Delta\boldsymbol{\pi}$); lower values indicate better estimation. We report averages and standard deviations over folds.

Execution Environment. For Tables 2 - 4, we measure timing on an Intel Xeon CPU E5-2680v2 2.8GHz with 128GB RAM. Particularly for experiments on larger synthetic datasets (c.f. Figures 1 - 2), we use an Intel Xeon CPU E5-2680v4 2.4GHz with 500GB RAM.

Sample CV. We begin with the experiments on sample CV, in which we partition samples \mathcal{N} into training and test sets. Table 2 shows the evaluations of all algorithms trained on a synthetic dataset with $n = 1000$ samples, $p = 100$ features, $M = 1000$ observations, and query size $|A_\ell| = 2$. PLADMM and PLADMM-log converge 4–27 times faster than other feature methods, i.e., Newton on $\boldsymbol{\beta}$ and SLSQP. Recall that in sample CV partitioning, training observations contain *only* training samples: test samples do not participate in any of the training observations. Thus, featureless methods ILSR, MM, and Newton on $\boldsymbol{\theta}$ are *no better than random predictors*, with 0.5 Top-1 Acc. and 0.0 KT. By regressing the Plackett-Luce scores from features, PLADMM and PLADMM-log significantly outperform the predictions of ILSR, MM, and Newton on $\boldsymbol{\theta}$, by 16% – 33% Top-1 Acc. and 16% – 30% KT.

Real datasets. We observe an equally significant speed gain on real datasets; Table 3 shows the evaluations on four real datasets partitioned w.r.t. sample CV. PLADMM and PLADMM-log are 3 – 18 times faster than Newton on $\boldsymbol{\beta}$ and SLSQP. This speed gain is fundamentally due to the smaller per iteration complexity of PLADMM and PLADMM-log (c.f. Section 4). For instance, compared to Newton on $\boldsymbol{\beta}$, PLADMM-log re-

²Our code is publicly available at <https://github.com/neu-spiral/FastAndAccurateRankingRegression>

Partitioning	Method	Training Metrics			Performance Metrics on the Test Set	
		Time (s) ↓	Iter. ↓	$\Delta\pi$ ↓	Top-1 Acc. ↑	KT ↑
Sample CV	PLADMM	0.237 ± 0.006	4 ± 0	0.717 ± 0.207	0.831 ± 0.119	0.609 ± 0.247
	PLADMM-log	1.428 ± 2.595	49 ± 79	0.845 ± 0.204	0.668 ± 0.159	0.335 ± 0.318
	ILSR (no \mathbf{X})	0.045 ± 0.002	2 ± 0	0.718 ± 0.207	0.5 ± 0.0	-1.0 ± 0.0
	MM (no \mathbf{X})	9.728 ± 0.487	500 ± 0	1.2 ± 0.1	0.5 ± 0.0	0.0 ± 0.0
	Newton on θ (no \mathbf{X})	4.537 ± 0.729	14 ± 3	1.236 ± 0.132	0.5 ± 0.0	-0.08 ± 0.272
	Newton on β	6.406 ± 2.104	14 ± 5	0.808 ± 0.462	0.844 ± 0.148	0.688 ± 0.296
	SLSQP	43.908 ± 24.469	229 ± 132	0.718 ± 0.206	0.796 ± 0.106	0.592 ± 0.211
Observation CV	PLADMM	0.48 ± 0.24	4 ± 0	0.717 ± 0.207	0.837 ± 0.037	0.569 ± 0.072
	PLADMM-log	1.58 ± 2.027	29 ± 14	0.883 ± 0.208	0.699 ± 0.066	0.398 ± 0.132
	ILSR (no \mathbf{X})	0.098 ± 0.056	2 ± 0	0.718 ± 0.208	0.708 ± 0.045	0.389 ± 0.088
	MM (no \mathbf{X})	11.302 ± 0.515	500 ± 0	0.864 ± 0.19	0.685 ± 0.037	0.354 ± 0.074
	Newton on θ (no \mathbf{X})	8.218 ± 1.782	14 ± 3	1.244 ± 0.121	0.506 ± 0.029	0.01 ± 0.05
	Newton on β	7.696 ± 2.35	14 ± 4	0.804 ± 0.463	0.871 ± 0.087	0.742 ± 0.173
	SLSQP	47.824 ± 28.585	219 ± 138	0.718 ± 0.206	0.819 ± 0.035	0.637 ± 0.07

Table 2: Evaluations on a synthetic dataset with $n = 1000$, $p = 100$, and $M = 1000$, partitioned w.r.t. sample CV and observation CV. We report the convergence time (Time), number of iterations until convergence (Iter), norm error in estimating true Plackett-Luce scores ($\Delta\pi$), top-1 accuracy on the test set (Top-1 Acc.), and Kendall-Tau correlation on the test set (KT). ILSR, MM, and Newton on θ do not use the features \mathbf{X} . Newton on β and SLSQP regress π from \mathbf{X} .




Figure 1: Convergence time (Time) and top-1 test accuracy (Top-1 Acc.) vs. n and p for PLADMM, PLADMM-log, ILSR, and Newton on β . Evaluations are on synthetic datasets containing $M = 250$ observations partitioned w.r.t. observation CV. Number of samples changes in $n \in \{50, 100, 1000, 10000, 100000\}$ when number of features is $p = 100$, and number of features changes in $p \in \{10, 100, 1000, 10000\}$ when number of samples is $n = 1000$.

Dataset	Method	Training Metrics		Performance Metrics on the Test Set	
		Time (s) ↓	Iter. ↓	Top-1 Acc. ↑	KT ↑
FAC	PLADMM	0.301 ± 0.048	4 ± 0	0.654 ± 0.237	0.307 ± 0.473
	PLADMM-log	0.298 ± 0.466	10 ± 15	0.685 ± 0.237	0.369 ± 0.474
	ILSR (no \mathbf{X})	0.059 ± 0.016	2 ± 0	0.5 ± 0.0	-1.0 ± 0.0
	MM (no \mathbf{X})	5.905 ± 0.282	500 ± 0	0.5 ± 0.0	0.0 ± 0.0
	Newton on θ (no \mathbf{X})	7.604 ± 0.805	18 ± 2	0.5 ± 0.0	-0.4 ± 0.49
	Newton on β	0.859 ± 0.077	6 ± 1	0.67 ± 0.17	0.34 ± 0.339
	SLSQP	14.332 ± 5.684	178 ± 67	0.675 ± 0.147	0.349 ± 0.293
ROP	PLADMM	1.708 ± 0.166	4 ± 0	0.783 ± 0.03	0.565 ± 0.06
	PLADMM-log	0.325 ± 0.028	1 ± 0	0.724 ± 0.105	0.448 ± 0.209
	ILSR (no \mathbf{X})	0.649 ± 0.053	2 ± 0	0.5 ± 0.0	-1.0 ± 0.0
	MM (no \mathbf{X})	0.001 ± 0.001	1 ± 0	0.5 ± 0.0	-1.0 ± 0.0
	Newton on θ (no \mathbf{X})	68.924 ± 5.521	8 ± 0	0.497 ± 0.012	-0.988 ± 0.036
	Newton on β	47.563 ± 8.342	2 ± 1	0.552 ± 0.048	0.103 ± 0.096
	SLSQP	4.823 ± 4.914	2 ± 1	0.769 ± 0.052	0.538 ± 0.104
Pairwise Sushi	PLADMM	0.046 ± 0.01	4 ± 0	0.451 ± 0.082	-0.09 ± 0.177
	PLADMM-log	0.141 ± 0.025	27 ± 13	0.532 ± 0.076	0.064 ± 0.152
	ILSR (no \mathbf{X})	0.014 ± 0.006	2 ± 0	0.5 ± 0.0	-1.0 ± 0.0
	MM (no \mathbf{X})	1.513 ± 0.587	352 ± 183	0.5 ± 0.0	-1.0 ± 0.0
	Newton on θ (no \mathbf{X})	1.282 ± 0.924	18 ± 9	0.5 ± 0.0	-0.666 ± 0.472
	Newton on β	0.21 ± 0.115	4 ± 2	0.665 ± 0.035	0.33 ± 0.069
	SLSQP	4.619 ± 6.321	168 ± 235	0.624 ± 0.065	0.248 ± 0.13
Triplet Sushi	PLADMM	0.091 ± 0.02	4 ± 0	0.358 ± 0.805	-0.333 ± 0.924
	PLADMM-log	0.556 ± 0.276	40 ± 25	0.393 ± 0.826	0.096 ± 1.069
	ILSR (no \mathbf{X})	0.033 ± 0.012	2 ± 0	0.334 ± 0.0	-0.047 ± 1.089
	MM (no \mathbf{X})	1.824 ± 0.701	267 ± 151	0.334 ± 0.0	0.0 ± 0.0
	Newton on θ (no \mathbf{X})	2.728 ± 1.475	13 ± 3	0.334 ± 0.0	-0.047 ± 1.089
	Newton on β	1.966 ± 3.158	10 ± 19	0.322 ± 0.802	-0.261 ± 0.956
	SLSQP	1.656 ± 1.793	20 ± 30	0.608 ± 0.826	0.358 ± 0.928

Table 3: Evaluations on real datasets partitioned w.r.t. sample CV (c.f. Sec. 5). We report the convergence time (Time), number of iterations until convergence (Iter), top-1 accuracy on the test set (Top-1 Acc.), and Kendall-Tau correlation on the test set (KT). ILSR, MM, and Newton on θ do not use the features \mathbf{X} . Newton on β and SLSQP regress π from \mathbf{X} .

quires about 2 times more iterations, but still converges 3 times faster than Newton on β on FAC. Moreover, while significantly decreasing the convergence time, PLADMM or PLADMM-log consistently attain similar prediction performance to Newton on β and SLSQP, except for Sushi, for which they perform slightly worse (by 13% – 20% Top-1 Acc.), though the convergence time dividends are striking in comparison (78% – 95%).

Aligned with the prediction performance on synthetic datasets, featureless methods ILSR, MM, and Newton on θ can only attain 0.5 Top-1 Acc. and 0.0 KT. By regressing the Plackett-Luce scores from features, PLADMM and PLADMM-log significantly outperform the predictions of ILSR, MM, and Newton on θ , by 3% – 31% Top-1 Acc. and 5% – 78% KT.

Observation CV. A sample can appear in both training and test observations in observation CV. Hence, featureless methods ILSR, MM, and Newton on θ should fare better than in sample CV. Nonetheless, in Table 2, as $n = 1000$ is larger than $p = 100$, there are more scores to learn than parameters. As a result, feature methods are advantageous for good predictions compared to featureless methods. Particularly, PLADMM and PLADMM-log outperform the predictions of ILSR, MM, and Newton on θ in observation CV on Table 2, by 13% Top-1 Acc. and 9% KT. The relative performance of feature vs. featureless methods is governed by the relationship among n , p , and M . We therefore explore the effect of n and p below; the effect of M is discussed in Appendix F.4. We do not include Newton on θ and MM in this analysis, as they are too slow.

Impact of p . To assess the impact of number of parameters, we fix $n = 1000$, $M = 250$, $|A_\ell| = 2$ and generate synthetic datasets with $p \in \{10, 100, 1000, 10000\}$. Fig. 1a shows the Time and Top-1 Acc. of PLADMM, PLADMM-log, ILSR, and Newton on β . As $M = 250$ observations are not enough to learn $n = 1000$ scores, PLADMM leads to significantly better Top-1 Acc. compared to ILSR. When $p = 10$, PLADMM and PLADMM-log outperform ILSR by 18% – 28% Top-1 Acc. Moreover, PLADMM and PLADMM-log are consistently faster than Newton on β , for all $p > 100$. Particularly, for $p = 10000$, PLADMM and PLADMM-log converge *42-579 times faster than Newton on β* . Interestingly, the convergence time of PLADMM-log can even decrease with increasing p . This is because the number of iterations until convergence decreases. While significantly decreasing the convergence time, PLADMM consistently attains better Top-1 Acc. than Newton on β , up to 8% for $p = 100$.

Impact of n . To assess the impact of number of samples, we fix $p = 100$, $M = 250$, $|A_\ell| = 2$ and generate synthetic datasets with $n \in \{50, 100, 1000, 10000, 100000\}$;

Fig. 1b shows evaluations on the resulting datasets. For $n > p = 100$, i.e., when there are more scores to learn than parameters, PLADMM leads to significantly better Top-1 Acc. compared to ILSR. Particularly, for $n = 100000$, PLADMM outperforms ILSR by 25% Top-1 Acc. This confirms that, especially when the number of observations M is not sufficient to learn n scores, exploiting the features associated with the samples is crucial in attaining good prediction performance. As expected, convergence time of Newton on β is not significantly affected by n . Despite this, PLADMM and PLADMM-log are faster than Newton on β for all $n < 1000$. Particularly, for $n = 50$, PLADMM and PLADMM-log converge *2-75 times faster than Newton on β* . While decreasing the convergence time, PLADMM consistently attains better Top-1 Acc. than Newton on β , up to 19% for $n = 100000$.

Real datasets. We include the evaluations on real datasets partitioned w.r.t. observation CV in the Appendix (c.f. Table 4). Performance agrees with observations above regarding the dependence on n and p . For datasets where $n > M > p$, e.g., FAC, PLADMM and PLADMM-log significantly outperform the predictions of ILSR, by 10% Top-1 Acc. and 25% KT. For datasets where M is much larger than n (c.f. Table 1), feature methods lead to similar prediction performance to each other and slightly lower performance than ILSR, MM, and Newton on θ . Overall, PLADMM and PLADMM-log consistently converge faster than Newton on β and SLSQP, by 3 – 27 times across all real datasets.

6 Conclusions

We solve the maximum likelihood estimation problem for the Plackett-Luce scores via ADMM. We show that the scores are equivalent to the stationary distribution of a Markov Chain and propose spectral algorithms, PLADMM and PLADMM-log. We model the Plackett-Luce scores as affine and logistic functions of features. Extending these to more complex models, particularly to deep neural networks, is an interesting open problem. Our approach has the potential of training a neural network over a linear penalty w.r.t. scores, where the latter are regressed efficiently via a spectral method over the *quadratic* pairwise ranking data. This can lead to significant improvements over training time, making an epoch linear rather than quadratic in sample size.

Acknowledgments

We are supported by NIH (R01EY019474), NSF (SCH-1622542 at MGH; SCH-1622536 at Northeastern; SCH-1622679 at OHSU), and by unrestricted departmental funding from Research to Prevent Blindness (OHSU).

Bibliography

- Agarwal, A., Patil, P., and Agarwal, S. (2018). Accelerated spectral ranking. In *International Conference on Machine Learning*, pages 70–79.
- Agarwal, S. (2016). On ranking and choice models. In *IJCAI*, pages 4050–4053.
- Ammar, A. and Shah, D. (2011). Ranking: Compare, don’t score. In *2011 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 776–783. IEEE.
- Ataer-Cansızoğlu, E. (2015). *Retinal image analytics: A complete framework from segmentation to diagnosis*. Northeastern University.
- Azari, H., Parks, D., and Xia, L. (2012). Random utility theory for social choice. In *Advances in Neural Information Processing Systems*, pages 126–134.
- Blanchet, J., Gallego, G., and Goyal, V. (2016). A markov chain approximation to choice modeling. *Operations Research*, 64(4):886–905.
- Boyd, S., Parikh, N., Chu, E., Peleato, B., and Eckstein, J. (2011). Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122.
- Boyd, S. and Vandenberghe, L. (2004). *Convex optimization*. Cambridge university press.
- Bradley, R. A. and Terry, M. E. (1952). Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345.
- Braverman, M. and Mossel, E. (2008). Noisy sorting without resampling. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 268–276. Society for Industrial and Applied Mathematics.
- Burges, C., Shaked, T., Renshaw, E., Lazier, A., Deeds, M., Hamilton, N., and Hullender, G. (2005). Learning to rank using gradient descent. In *Proceedings of the International Conference on Machine Learning*, pages 89–96. ACM.
- Caron, F. and Doucet, A. (2012). Efficient bayesian inference for generalized bradley–terry models. *Journal of Computational and Graphical Statistics*, 21(1):174–196.
- Cattelan, M. (2012). Models for paired comparison data: A review with emphasis on dependent data. *Statistical Science*, pages 412–433.
- Chang, H., Yu, F., Wang, J., Ashley, D., and Finkelstein, A. (2016). Automatic triage for a photo series. *ACM Transactions on Graphics (TOG)*, 35(4):148.
- Chartrand, R. and Wohlberg, B. (2013). A nonconvex ADMM algorithm for group sparsity with sparse groups. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 6009–6013.
- Chen, Y. and Suh, C. (2015). Spectral mle: Top-k rank aggregation from pairwise comparisons. In *International Conference on Machine Learning*, pages 371–380.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and Fei-Fei, L. (2009). Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255.
- Dubey, A., Naik, N., Parikh, D., Raskar, R., and Hidalgo, C. A. (2016). Deep learning the city: Quantifying urban perception at a global scale. In *European Conference on Computer Vision*, pages 196–212. Springer.
- Dwork, C., Kumar, R., Naor, M., and Sivakumar, D. (2001). Rank aggregation methods for the web. In *Proceedings of the 10th International Conference on World Wide Web*, pages 613–622. ACM.
- Dykstra, O. (1960). Rank analysis of incomplete block designs: A method of paired comparisons employing unequal repetitions on pairs. *Biometrics*, 16(2):176–188.
- Elo, A. E. (1978). *The rating of chessplayers, past and present*. Arco Pub.
- Fligner, M. A. and Verducci, J. S. (1993). *Probability models and statistical analyses for ranking data*, volume 80. Springer.
- Gallager, R. G. (2013). *Stochastic Processes: Theory for Applications*. Cambridge University Press.
- Guiver, J. and Snelson, E. (2009). Bayesian inference for plackett-luce ranking models. In *Proceedings of the 26th Annual International Conference on Machine Learning*, pages 377–384. ACM.
- Guo, K., Han, D., and Wu, T.-T. (2017). Convergence of alternating direction method for minimizing sum of two nonconvex functions with linear constraints. *International Journal of Computer Mathematics*, 94(8):1653–1669.
- Guo, Y., Tian, P., Kalpathy-Cramer, J., Ostmo, S., Campbell, J. P., Chiang, M. F., Erdogmus, D., Dy, J. G., and Ioannidis, S. (2018). Experimental design under the bradley-terry model. In *IJCAI*, pages 2198–2204.
- Hajek, B., Oh, S., and Xu, J. (2014). Minimax-optimal inference from partial rankings. In *Advances in Neural Information Processing Systems*, pages 1475–1483.
- Han, B. (2018). Dateline: Deep plackett-luce model with uncertainty measurements. *arXiv preprint arXiv:1812.05877*.

- Hong, M. (2018). A distributed, asynchronous, and incremental algorithm for nonconvex optimization: An ADMM approach. *IEEE Transactions on Control of Network Systems*, 5(3):935–945.
- Horn, R. A. and Johnson, C. R. (2012). *Matrix Analysis*. Cambridge university press.
- Hunter, D. R. (2004). MM algorithms for generalized bradley-terry models. *The Annals of Statistics*, 32(1):384–406.
- Jang, M., Kim, S., Suh, C., and Oh, S. (2017). Optimal sample complexity of m-wise data for top-k ranking. In *Advances in Neural Information Processing Systems*, pages 1686–1696.
- Joachims, T. (2002). Optimizing search engines using clickthrough data. In *Proceedings of the eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 133–142. ACM.
- Jolliffe, I. T. (1986). Principal component analysis and factor analysis. In *Principal Component Analysis*. Springer.
- Kamishima, T., Hamasaki, M., and Akaho, S. (2009). A simple transfer learning method and its application to personalization in collaborative tagging. In *ICDM*.
- Kendall, M. G. (1938). A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.
- Khetan, A. and Oh, S. (2016). Computational and statistical tradeoffs in learning to rank. In *Advances in Neural Information Processing Systems*, pages 739–747.
- Lei, Q., Zhong, K., and Dhillon, I. S. (2016). Coordinate-wise power method. In *Advances in Neural Information Processing Systems*, pages 2064–2072.
- Lu, T. and Boutilier, C. (2011). Learning mallows models with pairwise preferences. In *Proceedings of the 28th International Conference on Machine Learning (icml-11)*, pages 145–152.
- Mallows, C. L. (1957). Non-null ranking models. i. *Biometrika*, 44(1/2):114–130.
- Mao, C., Pananjady, A., and Wainwright, M. J. (2018). Breaking the $1/\sqrt{n}$ barrier: Faster rates for permutation-based models in polynomial time. In *Conference On Learning Theory*, pages 2037–2042.
- Mao, C., Weed, J., and Rigollet, P. (2017). Minimax rates and efficient algorithms for noisy sorting. *arXiv preprint arXiv:1710.10388*.
- Marden, J. I. (2014). *Analyzing and modeling rank data*. Chapman and Hall/CRC.
- Maystre, L. and Grossglauser, M. (2015). Fast and accurate inference of plackett-luce models. In *Advances in Neural Information Processing Systems*, pages 172–180.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior.
- McFadden, D. (2000). Disaggregate behavioral travel demand’s rum side. *Travel Behaviour Research*, pages 17–63.
- Negahban, S., Oh, S., and Shah, D. (2012). Iterative ranking from pair-wise comparisons. In *Advances in Neural Information Processing Systems*, pages 2474–2482.
- Negahban, S., Oh, S., Thekumparampil, K. K., and Xu, J. (2018). Learning from comparisons and choices. *The Journal of Machine Learning Research*, 19(1):1478–1572.
- Niranjan, U. and Rajkumar, A. (2017). Inductive pair-wise ranking: going beyond the $n \log(n)$ barrier. In *Thirty-First AAAI Conference on Artificial Intelligence*.
- Nocedal, J. and Wright, S. (2006). *Numerical optimization*. Springer Science & Business Media.
- Pahikkala, T., Tsivtsivadze, E., Airola, A., Järvinen, J., and Boberg, J. (2009). An efficient algorithm for learning to rank from preference graphs. *Machine Learning*, 75(1):129–165.
- Plackett, R. L. (1975). The analysis of permutations. *Applied Statistics*, pages 193–202.
- Ragain, S. and Ugander, J. (2016). Pairwise choice markov chains. In *Advances in Neural Information Processing Systems*, pages 3198–3206.
- Rajkumar, A. and Agarwal, S. (2014). A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In *International Conference on Machine Learning*, pages 118–126.
- Rajkumar, A. and Agarwal, S. (2016). When can we rank well from comparisons of $o(n \log(n))$ non-actively chosen pairs? In *Conference on Learning Theory*, pages 1376–1401.
- Ryzin, G. v. and Mahajan, S. (1999). On the relationship between inventory costs and variety benefits in retail assortments. *Management Science*, 45(11):1496–1509.
- Saha, A. and Rajkumar, A. (2018). Ranking with features: Algorithm and a graph theoretic analysis. *arXiv preprint arXiv:1808.03857*.
- Shah, N., Balakrishnan, S., Guntuboyina, A., and Wainwright, M. (2016a). Stochastically transitive models for pairwise comparisons: Statistical and computational issues. In *International Conference on Machine Learning*, pages 11–20.
- Shah, N. B., Balakrishnan, S., Bradley, J., Parekh, A., Ramchandran, K., and Wainwright, M. J. (2016b).

- Estimation from pairwise comparisons: Sharp minimax bounds with topology dependence. *The Journal of Machine Learning Research*, 17(1):2049–2095.
- Soufiani, H. A., Chen, W., Parkes, D. C., and Xia, L. (2013). Generalized method-of-moments for rank aggregation. In *Advances in Neural Information Processing Systems*, pages 2706–2714.
- Sun, W.-T., Chao, T.-H., Kuo, Y.-H., and Hsu, W. H. (2017). Photo filter recommendation by category-aware aesthetic learning. *IEEE Transactions on Multimedia*, 19(8):1870–1880.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. (2015). Going deeper with convolutions. In *Computer Vision and Pattern Recognition, 2015*.
- Thurstone, L. L. (1927). The method of paired comparisons for social values. *The Journal of Abnormal and Social Psychology*, 21(4):384.
- Tian, P., Guo, Y., Kalpathy-Cramer, J., Ostmo, S., Campbell, J. P., Chiang, M. F., Dy, J., Erdogmus, D., and Ioannidis, S. (2019). A severity score for retinopathy of prematurity. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 1809–1819. ACM.
- Vojnovic, M. and Yun, S. (2016). Parameter estimation for generalized thurstone choice models. In *International Conference on Machine Learning*, pages 498–506.
- Wang, Y., Yin, W., and Zeng, J. (2019). Global convergence of ADMM in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, 78(1):29–63.
- Wauthier, F., Jordan, M., and Jojic, N. (2013). Efficient ranking from pairwise comparisons. In *International Conference on Machine Learning*, pages 109–117.
- Yıldız, İ., Tian, P., Dy, J., Erdoğan, D., Brown, J., Kalpathy-Cramer, J., Ostmo, S., Campbell, J. P., Chiang, M. F., and Ioannidis, S. (2019). Classification and comparison via neural networks. *Neural Networks*.
- Zeng, J., Ouyang, S., Lau, T. T.-K., Lin, S., and Yao, Y. (2018). Global convergence in deep learning with variable splitting via the kurdyka-lojasiewicz property. *arXiv preprint arXiv:1803.00225*.