

Inductive Pairwise Ranking: Going Beyond the $n \log(n)$ Barrier

U. N. Niranjan*

Department of Computer Science
University of California Irvine
un.niranjan@uci.edu

Arun Rajkumar

Data Analytics Lab
Xerox Research Centre India
arun.rajkumar@xerox.com

Abstract

We study the problem of ranking a set of items from non-actively chosen pairwise preferences where each item has feature information with it. We propose and characterize a very broad class of preference matrices giving rise to the *Feature Low Rank* (FLR) model, which subsumes several models ranging from the classic Bradley–Terry–Luce (BTL) (Bradley and Terry 1952) and Thurstone (Thurstone 1927) models to the recently proposed blade-chest (Chen and Joachims 2016) and generic low-rank preference (Rajkumar and Agarwal 2016) models. We use the technique of matrix completion in the presence of side information to develop the Inductive Pairwise Ranking (IPR) algorithm that provably learns a good ranking under the FLR model, in a sample-efficient manner. In practice, through systematic synthetic simulations, we confirm our theoretical findings regarding improvements in the sample complexity due to the use of feature information. Moreover, on popular real-world preference learning datasets, with as less as 10% sampling of the pairwise comparisons, our method recovers a good ranking.

Introduction

Ranking from pairwise comparisons or *preferences* is an ubiquitous problem in machine learning, statistics and theoretical computer science. In the so-called *non-active* setting, one is given comparison results of m pairs pre-selected from among all pairs of n items where each pair is compared at least K times. Particularly, the learner does not get to choose which pairs are to be compared. The goal is then to estimate a suitable ordering of the items, using the observed comparison results, that conforms to the true ordering, assuming one exists, up to the desired error ϵ in a suitably defined error measure.

In practical ranking applications, we often have side information associated with the items that need to be ranked – such a scenario is referred to as the *inductive* setting. An advantage in this inductive learning setting is that, in addition to ranking a given set of items, one is also able to rank new unseen items that may introduced after parameter learning. Motivated by these factors, we wish to leverage the

available side information to compute an ordering more efficiently than existing techniques. This is relevant in many practical applications; for instance, in addition to using a minimal amount of customer preference data, (a) using food characteristics like nutrition, preparation method, etc could help in finding the top-rated dishes of a restaurant, (b) using car features like engine type, body type, etc could help elicit useful trends for automotive industry.

Our Contributions: To the best of our knowledge, our work is the first to derive a provable and efficient method for ranking in the non-active inductive setting. Our novelty and technical contributions can be summarized along the following axes:

1. *Model:* We generalize existing models so that we can incorporate (a) features, and (b) feature correlations associated with the items to be ranked. We show that our model subsumes many existing and popular ranking models.
2. *Algorithm:* Our algorithm uses two key subroutines namely, (a) noisy inductive matrix completion, and (b) approximate pairwise ranking algorithm (Copeland 1951).
3. *Guarantee:* We derive the guarantee that our algorithm obtains, with high probability, an ϵ -accurate recovery using $\Omega(\max(\log n/\epsilon^2, d^4 \log^3 n/\epsilon^3 n^2))$ independent pairwise comparisons chosen uniformly at random.
4. *Experiments:* We substantiate our theoretical results by demonstrating sample complexity gains on both synthetic and real-world experiments.

We would like to emphasize upfront that it is the sole focus of this paper to study the practically motivated regime of $d \ll n$ in detail. Furthermore, we note that our sample complexity results do not violate the standard $\Omega(n \log n)$ lower bounds for comparison-based sorting algorithms since we develop an algorithm that effectively ranks in the *feature space* rather than the *item space*.

Related Work and Background

We now give a brief overview of relevant work in ranking models followed by a brief background regarding tools from inductive matrix completion theory which will be crucial in proving our sample complexity bounds.

*Part of the work done while interning at Xerox Research Centre India.

Copyright © 2017, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Ranking Models : In the simplest terms, the ranking problem involves estimating the *best* ordering items according to some observed preferences. A early thread of ranking literature has its beginnings in economics involving choice models (Luce 1959); other related works in social choice theory include (Lu and Boutilier 2011) and (Caragiannis, Procaccia, and Shah 2013). A certain deterministic version of the ranking problem is also studied as the *sorting* problem which is central in theoretical computer science.

1. *Random Utility (RU) models:* Starting with the seminal work of (Bradley and Terry 1952), the Bradley-Terry-Luce (BTL) model has become a landmark model for ranking. In the vanilla version of this model, the probability that item i beats item j is given by $P_{ij} = \frac{w_i}{w_i + w_j}$ where $\mathbf{w} \in \mathbb{R}_+^n$ is the parameter vector to be estimated from data; the i^{th} entry in w denotes the *score* associated with item i . Thurstone (Thurstone 1927) model is also a well-known statistical model; here, $P_{ij} = \Phi(s_i - s_j)$ where Φ is the standard normal Cumulative Distribution Function (CDF) and $\mathbf{s} \in \mathbb{R}^n$ is the score vector. These classic models fall under the so-called *Random Utility (RU) Models* (Marschak and others 1959).
2. *Item Feature (IF) models:* Extending the BTL model, statistical models that utilize side information are presented in (Cattelan 2012). Recently, (Chen and Joachims 2016) presented the blade-chest ranking model which studied the stochastic intransitive setting. Their algorithm involves regularized maximum likelihood estimation for which tight sample complexity properties are not known. Despite the above works, to the best of our knowledge, there are no known models utilizing feature information while have provable sample-efficient algorithms for estimation and ranking.
3. *Low Rank (LR) models:* Recently, (Rajkumar and Agarwal 2016) – unifying classic models such as BTL and Thurstone models – defined a generic class of preference matrices which have low rank under transformations involving suitable *link functions*. Upon such a transformation, connections of the ranking problem to matrix completion theory become clear. Subsequently, they use well-known matrix completion results to derive sample complexity guarantees for ranking. However, their model does not utilize side information that may be available.

This list is by no means exhaustive; while there exist several other ranking methods (eg, ranking-SVM(Joachims 2002)), there are no known sample complexity guarantees associated with these.

Inductive Matrix Completion: The matrix completion task (Candès and Recht 2009) is to fill-in the missing entries of a partially observed matrix, which is possible efficiently under a low-rank assumption on the underlying matrix. Oftentimes, side information may be available which further makes this task potentially easier. This is the *Inductive Matrix Completion* (IMC) problem which is formally defined as the optimization problem, $\hat{\mathbf{Z}} = \arg \min_{\mathbf{Z}} \ell((\mathbf{A}^\top \mathbf{Z} \mathbf{B})_{ij}, M_{ij})$ where $\mathbf{A} \in \mathbb{R}^{d_1 \times n_1}$ and $\mathbf{B} \in \mathbb{R}^{d_2 \times n_2}$

are known feature matrices, $\mathbf{Z} \in \mathbb{R}^{d_1 \times d_2}$ is a rank- r unknown latent parameter matrix, $(i, j) \in \Xi \subseteq [n] \times [n]$ is the support set corresponding to the (uniformly sampled) observed entries and ℓ is any loss function, the squared loss being the most commonly chosen one. Once the estimate $\hat{\mathbf{Z}}$ is obtained using the training set indexed by Ξ , predictions may then be performed as $\hat{M}_{ij} = (\mathbf{A}^\top \hat{\mathbf{Z}} \mathbf{B})_{ij}$ for any $(i, j) \in \Xi^c$. The known solution techniques with recovery guarantees are:

1. *Non-convex algorithm (via alternating minimization):* This approach entails parameterizing $\mathbf{Z} = \mathbf{U}\mathbf{V}^\top$ and performing alternating projected least squares updates on \mathbf{U} and \mathbf{V} . The tightest known guarantee for this approach involves a sample complexity of $\Omega(d^2 r^3 \kappa^2 \log(d))$ and a convergence rate of $O(\log(1/\epsilon))$ (Jain and Dhillon 2013).
2. *Convex relaxation (via trace-norm formulation):* This approach entails relaxing the rank constraints to a nuclear norm penalty. Existence of a unique optimum can be shown with high probability (Xu, Jin, and Zhou 2013) and is characterized a sample complexity of $\Omega(dr \log(d) \log(n))$. Despite the non-smoothness, a sub-gradient descent algorithm provably converges with a rate of $O(1/\sqrt{\epsilon})$ (Ji and Ye 2009). Noisy features are handled in (Chiang, Hsieh, and Dhillon 2015).

Feature-aware Ranking

Notation and Preliminaries

General notation: Unless stated otherwise, we use lower-case letters for scalars, upper-case letters for universal constants, lower-case bold-face letters for vectors and upper-case bold-face letters for matrices; specifically, \mathbf{P} denotes a preference matrix. For any matrix $\mathbf{M} \in \mathbb{R}^{a \times b}$, let $\|\mathbf{M}\|_\infty = \max_{i,j} |M_{ij}|$, $\|\mathbf{M}\|_* = \sum_{i=1}^{\min\{a,b\}} \sigma_i(\mathbf{M})$ where $\sigma_i(\mathbf{M})$ are the singular values of \mathbf{M} and $\|\mathbf{M}\|_F = \sqrt{\sum_{i=1}^a \sum_{j=1}^b M_{ij}^2}$.

\mathbf{I} denotes the identity matrix whose dimensions would be implied from the context; similarly, depending on the context, $\mathbf{0}$ denotes a vector or matrix of zeros of the appropriate dimension. Next, let $P_{\min} = \min_{i \neq j} P_{ij}$ and $\Delta = \min_{i \neq j} |\psi(P_{ij}) - \psi(1/2)|$. Let Ξ be the support set of the observed entries of a matrix and let $m = |\Xi|$. Define projection of a matrix on the support set $B = \mathcal{R}_\Xi(A)$ as: $B_{ij} = A_{ij}$ if $(i, j) \in \Xi$ and $B_{ij} = 0$ if $(i, j) \notin \Xi$.

Items and features: Let n be the number of items to be ranked. Let \mathcal{S}_n denote the symmetric group on n items. Let each item have a d -dimensional feature vector associated with it, ie, $\mathbf{f}_i \in \mathbb{R}^d, \forall i \in [n]$; concatenating these, we obtain the feature matrix $\mathbf{F} = [\mathbf{f}_1, \dots, \mathbf{f}_n] \in \mathbb{R}^{d \times n}$.

Link functions: Any $\psi : [0, 1] \rightarrow \mathbb{R}$ which is a strictly increasing bijective function is a valid link function. For example, ψ could be the logit function, which is the inverse of the sigmoid function, defined as, $\psi(x) := \log\left(\frac{x}{1-x}\right)$ for $x \in [0, 1]$; another example is the probit function defined as $\psi(x) = \Phi^{-1}(x)$ where Φ is the standard normal CDF. When we apply the link function to a matrix, we mean that the transformation applied entry-wise.

Preference matrices: Let $\mathcal{P}_n := \{\mathbf{P} \in [0, 1]^{n \times n} | P_{ij} + P_{ji} = 1\}$ denote the set of all pairwise preference matrices over n items. Let the set of stochastic-transitive matrices be $\mathcal{P}_n^{ST} := \{\mathbf{P} \in \mathcal{P}_n | P_{ij} > 1/2, P_{jk} > 1/2 \implies P_{ik} > 1/2\}$ and the set of stochastic-intransitive matrices be $\mathcal{P}_n^{SI} := \{\mathbf{P} \in \mathcal{P}_n | P_{ij} > 1/2, P_{jk} > 1/2 \implies P_{ik} < 1/2\}$. Let \mathcal{P}_n^{RU} be the set of preference matrices associated with unary random utility models (which are described in the next section). Let $\mathcal{P}_n^{IF} := \{\mathbf{P} \in \mathcal{P}_n | P_{ij} = \psi^{-1}(\mathbf{w}^\top (\mathbf{f}_i - \mathbf{f}_j))\}$ for some $\mathbf{w} \in \mathbb{R}^d$.

Let $r \leq n$. Define the set of preference matrices having rank- r under the link function ψ as $\mathcal{P}_n(\psi, r) := \{\mathbf{P} \in \mathcal{P}_n | \text{rank}(\psi(\mathbf{P})) \leq r\}$. Next, define the set of preference matrices having rank- r under the link function ψ with the associated feature matrix $\mathbf{A} \in \mathbb{R}^{d \times n}$ as $\mathcal{P}_n(\psi, r, \mathbf{A}) := \{\mathbf{P} \in \mathcal{P}_n(\psi, r) | \psi(\mathbf{P}) = \mathbf{A}^\top \mathbf{L} \mathbf{A}\}$ where $\mathbf{L} \in \mathbb{R}^{d \times d}$ is an unknown rank- r latent matrix (which is a function of the parameters of the ranking model) and $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n] \in \mathbb{R}^{d \times n}$ is the known feature matrix whose i^{th} column is the feature vector corresponding to the i^{th} item. Let $\kappa = \sigma_{\min}(\mathbf{A})/\sigma_{\max}(\mathbf{A})$ be the inverse condition number of the feature matrix \mathbf{A} . Let $i \succ_{\mathbf{P}} j$ iff $P_{ij} > 1/2$. Denoting the indicator function by $\mathbb{1}$, we define the distance between a permutation $\sigma \in \mathcal{S}_n$ and a preference matrix $\mathbf{P} \in \mathcal{P}_n$ as:

$$\text{dist}(\sigma, \mathbf{P}) := \binom{n}{2}^{-1} \sum_{i < j} \mathbb{1}((i \succ_{\mathbf{P}} j) \wedge (\sigma(i) \succ \sigma(j))) + \binom{n}{2}^{-1} \sum_{i < j} \mathbb{1}((j \succ_{\mathbf{P}} i) \wedge (\sigma(j) \succ \sigma(i)))$$

Note that the above distance measure essentially counts the fraction of pairs on which σ and \mathbf{P} disagree, and can be thought of as a normalized 0 – 1 loss function.

Feature Low Rank Model

Random Utility (RU) models, arising in discrete choice theory, dating back to (Marschak and others 1959), characterize the probability of an item i beating item j , P_{ij} , using a prior on the (latent) score associated with those items, $w_i \in \mathbb{R}$ and $w_j \in \mathbb{R}$. The most popular pairwise ranking models including BTL and Thurstone models fit in this framework. In particular, it is well-known that if $w_i \sim \text{Gumbel}(0, 1)$, we obtain the BTL model; for completeness, we justify it below:

$$P_{ij} = \Pr(w_i > w_j) = \Pr(w_i - w_j > 0) \stackrel{\xi_1}{=} \frac{e^{-(w_i - w_j)}}{1 + e^{-(w_i - w_j)}}$$

where ξ_1 follows from the fact that the difference of two independent standard Gumbel distributed random variables follows the standard logistic distribution. Similarly, if $w_i \sim \mathcal{N}(0, 1)$, we obtain the Thurstone model. The underlying commonality in these models is the simple observation that the prior distribution is on the scores, which are unary terms. Notably, the recent result by (Rajkumar and Agarwal 2016) shows that under the inverse transformation of the CDF of the difference of the latent score variables, the preference probability matrix is low-rank for BTL and Thurstone models. Further, they extended this result to a broader class of

low-rank models in which the preference matrices are low-rank when the link function is set to be this inverse CDF.

One angle of motivation for this paper stems from the intuitive thought that the scores associated with an item i in RU models can be generalized to functions involving, not just unary terms but also, pairwise terms, ie, the score of item i with respect to item j is given by an energy function E_{ij} that has a bilinear form. From this point onwards, for simplicity, we detail the generalization of the RU models encompassing the BTL model, ie, we posit that E_{ij} has a standard Gumbel distribution and consequently, we choose the link function ψ to be the logit function. It is noteworthy that our results will hold under any link function for the corresponding prior.

We now propose the energy-based generative model, which we call *Feature Low Rank (FLR)* model, defined via the preference matrix specified as follows:

$$P_{ij} = \frac{e^{-E_{ij}}}{e^{-E_{ij}} + e^{-E_{ji}}} \quad (1)$$

Here, we define the energy function associated with the pair of items (i, j) to be of the form $E_{ij} := \mathbf{f}_i^\top \mathbf{w} + \mathbf{f}_i^\top \mathbf{W} \mathbf{f}_j$ where $\mathbf{w} \in \mathbb{R}^d$ and $\mathbf{W} \in \mathbb{R}^{d \times d}$ are the unknown latent parameters (vector and matrix parameters respectively) to be estimated, and \mathbf{f}_i and \mathbf{f}_j are the known feature vectors associated with items i and j respectively. It is clear from Equation (1) that a key advantage of the proposed model is the additional ability to incorporate side information in the form of feature vectors and feature correlations in a latent space described by \mathbf{W} . In matrix notation,

$$\begin{aligned} \psi(\mathbf{P}) &= (\mathbf{1} \mathbf{g}^\top + \mathbf{F}^\top \mathbf{W}^\top \mathbf{F}) - (\mathbf{g} \mathbf{1}^\top + \mathbf{F}^\top \mathbf{W} \mathbf{F}) \\ &= (\Sigma \mathbf{V}^\top)^\top \mathbf{L} (\Sigma \mathbf{V}^\top) \end{aligned} \quad (2)$$

where $\mathbf{g} := \mathbf{F}^\top \mathbf{w}$ is column vector in \mathbb{R}^n , $\mathbf{1} \in \mathbb{R}^n$ is the all-ones column vector, $\mathbf{F} = \mathbf{U} \Sigma \mathbf{V}^\top$ is the full SVD of \mathbf{F} (such that $\mathbf{U} \in \mathbb{R}^{d \times d}$, $\mathbf{V} \in \mathbb{R}^{n \times n}$ are orthonormal matrices with $\Sigma \in \mathbb{R}^{d \times n}$ as the $d \times d$ diagonal matrix of singular values padded with zeros) and $\mathbf{L} := \mathbf{U}^\top (\mathbf{1} \mathbf{w}^\top - \mathbf{w} \mathbf{1}^\top + \mathbf{W}^\top - \mathbf{W}) \mathbf{U}$ (such that $\Sigma_{ii}^{-1} = \sigma_i^{-1}$ and $\Sigma_{ij}^{-1} = 0$ if $i \neq j$). It is now clear that the sufficient condition for $\mathbf{P} \in \mathcal{P}_n(\psi, r, \Sigma \mathbf{V}^\top)$ is that $\text{rank}(\mathbf{L}) \leq r$. Now, we describe the generality of the FLR model in Equation (1) by showing that it subsumes many existing models and has much more expressiveness.

Proposition 1. *The LR model is a special case of the FLR model, ie, $\mathcal{P}_n(\psi, r) \subseteq \mathcal{P}_n(\psi, r, \mathbf{A})$.*

Corollary 1. *Let $\mathbf{F} = \mathbf{I}$ and ψ be the logit link function. From Proposition 1, it is easy to see the following special cases from Equation (1).*

1. Let $\mathbf{W} = \mathbf{x} \mathbf{y}^\top$. If $\mathbf{w} = 0$, then $\mathbf{P} \in \mathcal{P}_n(\psi, 2)$. If $\mathbf{w} \neq 0$, then $\mathbf{P} \in \mathcal{P}_n(\psi, 4)$.
2. If \mathbf{W} is symmetric, then $\mathbf{W} - \mathbf{W}^\top = 0$ and hence $\mathbf{P} \in \mathcal{P}_n(\psi, 2)$.
3. Let \mathbf{A} be a diagonal $r \times r$ matrix; let $\{\mathbf{X}, \mathbf{Y}\} \in \mathbb{R}^{n \times r}$ be orthonormal matrices. If $\mathbf{W} = \mathbf{X} \mathbf{A}_{r \times r} \mathbf{Y}^\top + \mathbf{M}$ where \mathbf{M} is a symmetric matrix, then $\psi(\mathbf{P}) \in \mathcal{P}_n(\psi, 2r + 2)$.

Table 1: Comparison of this work to previous works in non-active pairwise ranking: extended-BTL model is due to (Cattelan 2012), chest-blade model is due to (Chen and Joachims 2016) and low-rank model is due to (Rajkumar and Agarwal 2016). Here, \mathbf{w} and \mathbf{W} are model parameters as given in Equation (1), and \mathbf{F} denotes the feature matrix. Furthermore, *pairs* and *per* denote the state-of-the-art bounds known regarding the number of pairs compared and the number of comparisons per pair respectively. Note that we consider the practically important regime of $n \gg d$.

Model	\mathbf{F}	\mathbf{w}	\mathbf{W}	<i>pairs</i>	<i>per</i>
BTL	$\mathbf{F} = \mathbf{I}$	$\mathbf{w} \in \mathbb{R}^n$	$\mathbf{W} = \mathbf{0}$	$\Omega(n \log n)$	$\Omega(\log n)$
Item-feature	$\mathbf{F} \in \mathbb{R}^{d \times n}$	$\mathbf{w} \in \mathbb{R}^d$	$\mathbf{W} = \mathbf{0}$	many	many
Chest-blade	$\mathbf{F} = \mathbf{I}$	$\mathbf{w} \in \mathbb{R}^n$	$\text{rank}(\mathbf{W}) = O(d)$	many	many
Low-rank	$\mathbf{F} = \mathbf{I}$	$\mathbf{w} \in \mathbb{R}^n$	$\mathbf{W} \in \mathbb{R}^{n \times n}$	$\Omega(nr \log n)$	$\Omega(r \log n)$
This work	$\mathbf{F} \in \mathbb{R}^{d \times n}$	$\mathbf{w} \in \mathbb{R}^d$	$\mathbf{W} \in \mathbb{R}^{d \times d}$	$\Omega(d^2 \log n)$	$\Omega(d^2 \log^2 n / n^2)$

Proposition 2. *The unary RU models are special cases of the FLR model, ie, $\mathcal{P}_n^{\text{RU}} \subseteq \mathcal{P}_n(\psi, r, \mathbf{A})$.*

Corollary 2. *The BTL and Thurstone models are obtained as special cases of the FLR model under the logit and the probit transformations of \mathbf{P} respectively. This follows from Proposition 1 (or Corollary 1-part (1)) above together with Propositions 6 and 7 of (Rajkumar and Agarwal 2016).*

Proposition 3. *Regression-based models with item-specific features in (Cattelan 2012) are special cases of the FLR model, ie, $\mathcal{P}_n^{\text{IF}} \subseteq \mathcal{P}_n(\psi, r, \mathbf{A})$.*

Corollary 3. *Let $d \ll n$. Then we recover the blade-chest model (Chen and Joachims 2016) as a special case of the FLR model by setting $\text{rank}(\mathbf{W}) = O(d)$ and $\mathbf{w} = \mathbf{0}$. Next, when $d \geq n$, it is clear from Theorem 1 of (Chen and Joachims 2016) that such preference matrices degenerate into matrices in $\mathcal{P}_n(\psi, n, \mathbf{A})$ where ψ is the logit function. Moreover, it is easy to see that the FLR model admits both stochastic-transitive and stochastic-intransitive preference matrices.*

Due to space constraints, proofs of Propositions 1, 2 and 3 are given in the appendix. To summarize, we have shown how to instantiate several previously proposed ranking models as special cases of our FLR model in Table 1.

Problem Setup and Solution Approach

Once we have the generative ranking model as developed in the previous section, the objective in our learning problem is then to find the permutation of n items that minimizes the number of violations with respect to the true underlying preference matrix \mathbf{P} , ie, to find the best ranking $\hat{\sigma}$ in the sense that,

$$\hat{\sigma} = \arg \min_{\sigma} \text{dist}(\sigma, \mathbf{P})$$

The input is the pairwise comparison dataset $S = \{(i, j, y_{ij}^k)\}$ which consists of comparison results of pairs (i, j) from a survey involving K users where each user with index k assigns $y_{ij}^k = 1$ if he prefers i to j and $y_{ij}^k = 0$ if he prefers j to i . Note that it is not necessary that all pairs of items be compared; our algorithm is able to handle noisy and incomplete data. Since the true preference matrix \mathbf{P} is unknown, our algorithm instead proceeds by using the empirical preference matrix $\hat{\mathbf{P}}$ computed from the available y_{ij}^k ; it is to be noted, even then, our analysis guarantees that

$\text{dist}(\hat{\sigma}, \mathbf{P})$ is good as opposed to just $\text{dist}(\hat{\sigma}, \hat{\mathbf{P}})$. Additionally, in our inductive setting, the feature information is encoded by $\mathbf{f}_i \in \mathbb{R}^d$ for every item i and concatenated to form the feature matrix $\mathbf{F} \in \mathbb{R}^{d \times n}$.

Algorithm

We present our main algorithm for inductive ranking in Algorithm 3. The input data consist of the set of pairwise comparison results $S = \{(i, j, \{y_{ij}^k\})\}$, $(i, j) \in \Xi \subseteq [n] \times [n]$, $k \in [K]$, $y_{ij}^k \in \{0, 1\}$ and the feature matrix $\mathbf{F} \in \mathbb{R}^{d \times n}$. The algorithm assumes the link function and the rank as input parameters. The subroutines used are:

1. *Noisy matrix completion with features (Subroutine 1):* Note that to solve our ranking problem and derive the associated recovery guarantee, it suffices, as we have done, to use the specified trace-norm program as a black-box method; hence, we assume that we have access to an oracle that gives us the solution to the convex program. The details of how the solution to this program may be found numerically is beyond the scope of this work – for further details regarding some possible sub-gradient algorithms, we refer the reader to (Chiang, Hsieh, and Dhillon 2015) and (Ji and Ye 2009).
2. *γ -approximate pairwise ranking procedure (Subroutine 2):* Let $\hat{\sigma} \in \mathcal{S}_n$ be the output of any Pairwise Ranking (PR) procedure with respect to an underlying preference matrix \mathbf{P} . For a constant $\gamma > 1$, $\hat{\sigma}$ is said to be γ -approximate if $\text{dist}(\hat{\sigma}, \mathbf{P}) \leq \gamma \min_{\sigma \in \mathcal{S}_n} \text{dist}(\sigma, \mathbf{P})$. Any constant factor approximate ranking procedure maybe used. Specifically, we use the Copeland procedure (Copeland 1951) as a black-box method which has a 5-approximation guarantee (Coppersmith, Fleischer, and Rudra 2006). This method involves simply sorting the items according to a score which is computed for every item i as $\sum_{j=1}^n \mathbb{1}(\hat{P}_{ij} > 1/2)$.

Analysis

In this section, we state and prove our main result.

Theorem 1 (Guaranteed rank aggregation with sub-linear sample complexity using item features). *Let $\mathbf{P} \in (\mathcal{P}_n(\psi, r, \mathbf{A}) \cap \mathcal{P}_n^{\text{ST}})$ be the true underlying preference matrix according to which the pairwise comparison dataset $S = \{(i, j, \{y_{ij}^k\})\}$ is generated. Let ψ be L -Lipschitz*

Subroutine 1 IMC: Inductive Matrix Completion

Input: M_{ij} for $(i, j) \in \Xi \subseteq [n] \times [n]$, feature matrix \mathbf{F} .**Output:** Completed matrix $\bar{\mathbf{M}}$.

1: Solve the convex program:

$$\hat{\mathbf{Z}} = \arg \min_{\mathbf{Z}_L} \|\mathcal{R}_\Xi(\mathbf{M} - \mathbf{F}^\top \mathbf{Z}_L \mathbf{F})\|_F^2 \text{ s.t. } \|\mathbf{Z}_L\|_* \leq C_L$$

2: **return** $\psi(\hat{\mathbf{P}}) \leftarrow \mathbf{F}^\top \hat{\mathbf{Z}} \mathbf{F}$.

Subroutine 2 PR: Pairwise Ranking (Copeland Procedure)

Input: Preference matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$.**Output:** Ranking $\hat{\sigma}$.1: Threshold: $\forall (i, j), \quad \widetilde{M}_{ij} \leftarrow \mathbb{1}(M_{ij} > 1/2)$.2: Compute row-sum of $\widetilde{\mathbf{M}}$: $\mathbf{v} \leftarrow \widetilde{\mathbf{M}} \mathbf{1}$.3: **return** $\hat{\sigma} \leftarrow \text{Sort}(\mathbf{v})$.

in $[\frac{P_{\min}}{2}, 1 - \frac{P_{\min}}{2}]$. Let Ξ be the set of pairs of items compared such that the number of pairs compared is $|\Xi| = m > \frac{48C_2^2 d^2 \log(n)(1+\gamma)^2}{\kappa^8 \epsilon^2 \Delta^4}$ where Ξ is chosen uniformly at random from among all possible subsets of item pairs of size m . Let each pair in Ξ be compared independently $K \geq \frac{16(1+\gamma)mL^2 \log(n)}{n^2 \Delta^2 \epsilon}$ times where $\Delta = \min_{i \neq j} |\psi(P_{ij}) - \psi(1/2)|$. Then, with probability at least $1 - 3/n^3$, for any $\epsilon > 0$, Algorithm 3 returns an estimated permutation $\hat{\sigma}$ such that $\text{dist}(\hat{\sigma}, \mathbf{P}) \leq \epsilon$.

Remark 1. The key take-away message in Theorem 1 is the reduction in sample complexity possible due to efficient utilization of features and feature correlations, associated with the items to be ranked, by Algorithm 3. For instance, when $d = O(1)$, which is often the case in practice, we reduce the required total number of comparisons to be made to $\Omega(\log(n))$. Thus, we achieve a very significant gain since the total number of comparisons is poly-logarithmic as opposed to quadratic in the number of items. This is especially crucial in large-scale machine learning applications.

Remark 2. Another point to be noted from Theorem 1 is that, under the uniform sampling assumption, when features associated with items are known, it is more important that we compare sufficient (precisely, $\Omega(\log n)$) number of different pairs rather than high number of comparisons per pair. Furthermore, the total number of comparisons needed in Theorem 1 is given by the product mK which is $\Omega(\max(\log n/\epsilon^2, d^4 \log^3 n/\epsilon^3 n^2))$.

We now present the proof of Theorem 1. We shall prove the theorem under the Bernoulli sampling model (where each entry of an $n \times n$ matrix is observed independently with probability $1/n^2$) rather than the uniform sampling model (wherein Ξ is chosen uniformly at random from among all possible subsets of item pairs of size m); the equivalence between the two is well-known (see, for instance, Section 7.1 of (Candès et al. 2011)).

Proof. Let \hat{P}_{ij} be the empirical probability estimate of P_{ij} . Note that we compute $\hat{P}_{ij} = \frac{1}{K} \sum_{k=1}^K y_{ij}^k$ for $(i, j) \in$

Algorithm 3 IPR: Inductive Pairwise Ranking

Input: Set of comparison results $S = \{(i, j, \{y_{ij}^k\})\}$, feature matrix \mathbf{F} , link function ψ , target rank r .**Output:** Ranking of n items, $\hat{\sigma} \in \mathcal{S}_n$.1: Construct the partially observed empirical preference matrix using S as:

$$\hat{P}_{ij} = \begin{cases} \frac{1}{K} \sum_{k=1}^K y_{ij}^k & \text{if } (i, j) \in \Xi \\ \frac{1}{K} \sum_{k=1}^K (1 - y_{ij}^k) & \text{if } (j, i) \in \Xi \\ 1/2 & \text{if } i = j \text{ or } (i, j) \notin \Xi \end{cases}$$

2: Compute SVD of $\mathbf{F} = \mathbf{U}\Sigma\mathbf{V}^\top$ and set $\mathbf{A} \leftarrow \Sigma\mathbf{V}^\top$.

3: Use a noisy inductive matrix completion subroutine:

$$\psi(\hat{\mathbf{P}}) \leftarrow \text{IMC}(\psi(\hat{\mathbf{P}}), \mathbf{A}).$$

4: Take the inverse transform of the truncated r -SVD of the completed matrix estimate: $\mathbf{Q} \leftarrow \psi^{-1}(\mathcal{P}_r(\psi(\hat{\mathbf{P}})))$.5: Using a pairwise ranking subroutine: $\hat{\sigma} \leftarrow \text{PR}(\mathbf{Q})$.6: **return** $\hat{\sigma}$.

Ξ from the given pairwise comparison dataset, $S = \{(i, j, \{y_{ij}^k\})\}$. From Equation (2), $\psi(\mathbf{P}) = \mathbf{A}^\top \mathbf{L} \mathbf{A}$ where $\mathbf{A} = \Sigma\mathbf{V}^\top$. Since we use the empirical estimate for P_{ij} , we have noise due to sampling error only over Ξ , ie, $\psi(\hat{\mathbf{P}}) = \psi(\mathbf{P}) + \mathbf{N} = \mathbf{A}^\top \mathbf{L} \mathbf{A} + \mathbf{N}$ where

$$|N_{ij}| = \begin{cases} 0 & \text{if } (i, j) \notin \Xi \\ \left| \psi\left(\frac{1}{K} \sum_{k=1}^K y_{ij}^k\right) - \psi(P_{ij}) \right| & \text{if } (i, j) \in \Xi \end{cases}$$

Now, we solve the trace-norm regularized convex program corresponding to the noisy inductive matrix completion problem:

$$\{\bar{\mathbf{L}}, \bar{\mathbf{N}}\} = \arg \min_{\mathbf{Z}_N, \mathbf{Z}_L} \left\| \mathcal{R}_\Xi(\psi(\hat{\mathbf{P}}) - (\mathbf{A}^\top \mathbf{Z}_L \mathbf{A} + \mathbf{Z}_N)) \right\|_F^2 + \lambda_L \|\mathbf{Z}_L\|_* + \lambda_N \|\mathbf{Z}_N\|_*$$

and let $\psi(\hat{\mathbf{P}}) = \mathbf{A}^\top \bar{\mathbf{L}} \mathbf{A} + \bar{\mathbf{N}}$ be the link-transformed completed (estimate) matrix where $\bar{\mathbf{N}}$ be the estimated noise matrix. This is equivalent to solving the problem:

$$\{\bar{\mathbf{L}}, \bar{\mathbf{N}}\} = \arg \min_{\mathbf{Z}_N, \mathbf{Z}_L} \left\| \mathcal{R}_\Xi(\psi(\hat{\mathbf{P}}) - (\mathbf{A}^\top \mathbf{Z}_L \mathbf{A} + \mathbf{Z}_N)) \right\|_F^2 \text{ s.t. } \|\mathbf{Z}_L\|_* \leq C_L, \quad \|\mathbf{Z}_N\|_* \leq C_N$$

We set $C_N = 0$ and $C_L = \left\| (\mathbf{A}^\top)^\dagger \psi(\hat{\mathbf{P}}) (\mathbf{A})^\dagger \right\|_*$ which may be upper bounded, by Lemma 3 of (Chiang, Hsieh, and Dhillon 2015) as $C_L \leq \frac{d}{C' \kappa^4}$ for a constant C' . We now recall Theorem 1 from (Chiang, Hsieh, and Dhillon 2015). Let $\delta < 1/d$ and \mathbf{A} be well-conditioned, specifically, $\kappa^4 \leq C_2 d$ for constant C_2 . The expected squared loss under Bernoulli sampling is bounded as, with probability at least $1 - \delta$:

$$\frac{\left\| \psi(\hat{\mathbf{P}}) - \psi(\mathbf{P}) \right\|_F^2}{n^2} \leq C_1 \min \left(C_N \sqrt{\frac{\log(2n)}{m}}, \sqrt{C_N \frac{\sqrt{n}}{m}} \right) + \frac{C_2 d}{\kappa^4} \sqrt{\frac{\log(2/\delta)}{m}} \quad (3)$$

Figure 1: Ranking results of LRPR and IPR: fixing $d = 20$ and $K = 50\lceil d^2 \log^2(n)/n^2 \rceil$ while varying m .





Figure 2: Ranking results of LRPR and IPR: fixing $d = 20$ and $m = \lceil d^2 \log(n) \rceil$ while varying K .



where C_1 and C_2 are constants. By triangle inequality,

$$\begin{aligned} \|\psi(\bar{\mathbf{P}}) - \psi(\hat{\mathbf{P}})\|_F &= \|\psi(\bar{\mathbf{P}}) - (\psi(\mathbf{P}) + \mathbf{N})\|_F \\ &\geq \|\psi(\bar{\mathbf{P}}) - \psi(\mathbf{P})\|_F - \|\mathbf{N}\|_F \end{aligned}$$

Using $C_N = 0$ in Equation (3), with probability at least $1 - \delta$,

$$\frac{1}{n} \|\psi(\bar{\mathbf{P}}) - \psi(\mathbf{P})\|_F \leq \left(\frac{C_2 d}{\kappa^4} \sqrt{\frac{\log(2/\delta)}{m}} \right)^{1/2} + \frac{1}{n} \|\mathbf{N}\|_F$$

Let $K \geq \frac{m L^2 \log(n)}{\tau^2}$ where $\tau = n \sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta}{4}$. Substituting the bounds for the \mathbf{N} terms from Lemma 1 and using the union bound, with probability at least $1 - \delta - 1/n^3$,

$$\begin{aligned} \|\psi(\bar{\mathbf{P}}) - \psi(\mathbf{P})\|_F &\leq n \left(\frac{C_2 d}{\kappa^4} \sqrt{\frac{\log(2/\delta)}{m}} \right)^{\frac{1}{2}} + \tau \\ &\leq n \left(\frac{C_2 d}{\kappa^4} \sqrt{\frac{\log(2/\delta)}{m}} \right)^{\frac{1}{2}} + n \sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta}{4} \end{aligned}$$

Now, setting $m > \frac{16 C_2^2 d^2 \log(2/\delta) (1+\gamma)^2}{\kappa^8 \epsilon^2 \Delta^4}$ and $\delta = 2/n^3$, we obtain, with probability $1 - 3/n^3$, $\|\psi(\bar{\mathbf{P}}) - \psi(\mathbf{P})\|_F \leq n \sqrt{\frac{\epsilon}{1+\gamma}} \frac{\Delta}{2}$. Then, arguments similar to the proof of Theorem 13 of (Rajkumar and Agarwal 2016) yield our result. \square

Lemma 1 (Characterization of noise due to finite-sample effects). *Under the conditions of Theorem 1, let m item pairs be compared such that the number of comparisons per item pair is $K \geq \frac{m L^2 \log(n)}{\tau^2}$. Then, with probability at least $1 - 1/n^3$, $\|\mathbf{N}\|_F \leq \tau$.*

Due to space limitations, the proof of Lemma 1 is given in the appendix.

Figure 3: Ranking results of LRPR and IPR: fixing $K = 50\lceil d^2 \log^2(n)/n^2 \rceil$ and $m = \lceil d^2 \log(n) \rceil$ while varying d .





Figure 4: Ranking results of LRPR and IPR on Sushi and Car datasets: we fix $m = \lceil d^2 \log(n) \rceil$ while varying K ; we fix $K = 50\lceil d^2 \log^2(n)/n^2 \rceil$ while varying m .



Experimental Results

In this section, we conduct a systematic empirical investigation of the performance of our ranking method and justify our theoretical claim in the previous section. The goal of this study is two-fold: (a) to verify the correctness of our algorithm, and (b) to show that by using features and feature correlations, our IPR algorithm has a better sample complexity thereby improving upon the LRPR algorithm that does not take into account the available side information.

Synthetic Simulations

For a given set of $n = 500$ items, we consider three main problem parameters: (1) m – the number of item pairs compared (Figure 1), (2) K – the number of comparisons per pair (Figure 2), (3) d – the dimensionality of features (Figure 3). We study the performance of both IPR and LRPR algorithms by varying each of the problem parameters while fixing the others. We note that by making use of side information, IPR outperforms LRPR in all the cases as shown in the sample complexity plots. All the accuracy results presented are obtained by averaging over five runs.

Data generation: We consider three representative preference matrices derived from Equation (1): (a) *Model-1*: we set $W = 0$, (b) *Model-2*: we construct a general W ; here, we generate $W_{ij} \sim U(0, 1)$, and (c) *Model-3*: we construct a low-rank \mathbf{W} , ie, $\text{rank}(\mathbf{W}) = 2 < d$; here we generate $W_{ij} \sim U(0, 1)$ and then truncating \mathbf{W} by setting all but its top two singular values to zero. In all the three models, we generate $w_i \sim U(0, 1)$. The features are generated as $F_{ij} \sim U(0, 1)$; to ensure that the features are well-conditioned, we perform the full SVD of feature matrix \mathbf{F} and set all its singular values to 1.

Parameter settings: For IPR, we choose $\lambda_L = 10^{-2}$ and $\lambda_N = 10^2$. Note that LRPR allows for the rank of the problem to be automatically determined. In the same spirit, though Step-5 of Algorithm 3 requires the knowledge of

the true rank, we choose *not* to perform this truncation step thereby including the error induced by the smaller singular values resulting from noise due to sampling in our distance estimate – even then, IPR outperforms LRPR.

Real-data Simulations

We apply our method on two popular preference learning datasets. We briefly describe the data and the results (Figure 4) we obtain below:

1. *Sushi*: This data (Kamishima and Akaho 2009) is from a survey of 5000 customers. Each customer orders 10 sushi dishes according to their preferences. The goal, then, is to estimate a global ranking of these sushi dishes using these observations from customers. Each sushi has six features such as price, taste and so on. We construct the complete preference matrix $\mathbf{P} \in [0, 1]^{10 \times 10}$ using the preferences of all the customers and consider this to be ground truth preference matrix. An interesting observation was that, over five runs of the algorithms, IPR gets two out of the top four sushi dishes right most of the times namely, ‘amaebi’ and ‘ikura’; on the other hand, LRPR does not succeed in recovering these always.
2. *Car*: The task in this dataset (Abbasnejad et al. 2013) is find an order of preference among ten cars. This data was collected by surveying 60 customers regarding their preferences among pairs of cars drawn from the set of ten cars. Each car has four features including engine, transmission and so on. We construct the ground truth preference matrix $\mathbf{P} \in [0, 1]^{10 \times 10}$ by aggregating the pairwise preferences of all the customers. An interesting trend we found was that customers generally preferred sedans over SUVs and non-hybrid vehicles over hybrid vehicles.

Discussion and Future Directions

In this paper, we have proposed and characterized the FLR model together with the guaranteed IPR algorithm that utilizes available side information of the items to be ranked to provably reduce the sample complexity for ranking from $\Omega(n \log n)$ to possibly as low as $\Omega(\log n)$. A future research direction is to see if mixture models for ranking such as the recently proposed topic modeling approach (Ding, Ishwar, and Saligrama 2015) could fit into our framework while admitting sample-efficient estimation algorithms.

References

- Abbasnejad, E.; Sanner, S.; Bonilla, E. V.; Poupart, P.; et al. 2013. Learning community-based preferences via dirichlet process mixtures of gaussian processes. In *IJCAI*.
- 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.
- Candès, E. J., and Recht, B. 2009. Exact matrix completion via convex optimization. *Foundations of Computational mathematics* 9(6):717–772.
- Candès, E. J.; Li, X.; Ma, Y.; and Wright, J. 2011. Robust principal component analysis? *Journal of the ACM (JACM)* 58(3):11.
- Caragiannis, I.; Procaccia, A. D.; and Shah, N. 2013. When do noisy votes reveal the truth? In *Proceedings of the fourteenth ACM conference on Electronic commerce*, 143–160. ACM.
- Cattelan, M. 2012. Models for paired comparison data: A review with emphasis on dependent data. *Statistical Science* 412–433.
- Chen, S., and Joachims, T. 2016. Modeling intransitivity in matchup and comparison data. In *Proceedings of the Ninth ACM International Conference on Web Search and Data Mining*, 227–236. ACM.
- Chiang, K.-Y.; Hsieh, C.-J.; and Dhillon, I. S. 2015. Matrix completion with noisy side information. In *Advances in Neural Information Processing Systems*, 3447–3455.
- Copeland, A. H. 1951. A reasonable social welfare function. In *Mimeographed notes from a Seminar on Applications of Mathematics to the Social Sciences*, University of Michigan.
- Coppersmith, D.; Fleischer, L.; and Rudra, A. 2006. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, 776–782. Society for Industrial and Applied Mathematics.
- Ding, W.; Ishwar, P.; and Saligrama, V. 2015. A topic modeling approach to ranking. In *AISTATS*.
- Jain, P., and Dhillon, I. S. 2013. Provable inductive matrix completion. *arXiv preprint arXiv:1306.0626*.
- Ji, S., and Ye, J. 2009. An accelerated gradient method for trace norm minimization. In *Proceedings of the 26th annual international conference on machine learning*, 457–464. ACM.
- Joachims, T. 2002. Optimizing search engines using click-through data. In *Proceedings of the eighth ACM SIGKDD international conference on Knowledge discovery and data mining*, 133–142. ACM.
- Kamishima, T., and Akaho, S. 2009. Efficient clustering for orders. In *Mining complex data*. Springer. 261–279.
- Lu, T., and Boutilier, C. 2011. Learning mallows models with pairwise preferences. In *Proceedings of the 28th international conference on machine learning (icml-11)*, 145–152.
- Luce, R. D. 1959. *Individual Choice Behavior a Theoretical Analysis*. John Wiley and sons.
- Marschak, J., et al. 1959. Binary choice constraints on random utility indicators. Technical report, Cowles Foundation for Research in Economics, Yale University.
- Rajkumar, A., and Agarwal, S. 2016. When can we rank well from comparisons of $o(n \log(n))$ non-actively chosen pairs? In *29th Annual Conference on Learning Theory*, 1376–1401.
- Thurstone, L. L. 1927. A law of comparative judgment. *Psychological review* 34(4):273.
- Xu, M.; Jin, R.; and Zhou, Z.-H. 2013. Speedup matrix completion with side information: Application to multi-label learning. In *Advances in Neural Information Processing Systems*, 2301–2309.