

Fast and Robust Rank Aggregation against Model Misspecification

Yuangang Pan*

YUANGANG.PAN@GMAIL.COM

*Center for Frontier AI Research
Research Agency for Science, Technology and Research (A*STAR)
Singapore*

*and
Australian Artificial Intelligence Institute
University of Technology Sydney
NSW 2007, Australia*

Ivor W. Tsang[†]

IVOR.TSANG@GMAIL.COM

*Center for Frontier AI Research
Research Agency for Science, Technology and Research (A*STAR)
Singapore*

*and
Australian Artificial Intelligence Institute
University of Technology Sydney
NSW 2007, Australia*

Weijie Chen

WJCPER2008@126.COM

*Zhijiang College
Zhejiang University of Technology
Hangzhou 310014, Zhejiang, China*

Gang Niu

GANG.NIU@RIKEN.JP

*Center for Advanced Intelligence Project
RIKEN, Tokyo, 103-0027, Japan*

Masashi Sugiyama

SUGI@K.U-TOKYO.AC.JP

*Center for Advanced Intelligence Project
RIKEN, Tokyo, 103-0027, Japan*
*and
Graduate School of Frontier Sciences
University of Tokyo
Chiba 277-8561, Japan*

Editor: Sathiya Keerthi

Abstract

In rank aggregation (RA), a collection of preferences from different users are summarized into a total order under the assumption of homogeneity of users. Model misspecification in RA arises since the homogeneity assumption fails to be satisfied in the complex real-world situation. Existing robust RAs usually resort to an augmentation of the ranking model to account for additional noises, where the collected preferences can be treated as a noisy perturbation of idealized preferences. Since the majority of robust RAs rely on certain perturbation assumptions, they cannot generalize well to

*. Preliminary work was done during an internship at RIKEN AIP.

†. The corresponding author.

agnostic noise-corrupted preferences in the real world. In this paper, we propose CoarsenRank, which possesses robustness against model misspecification. Specifically, the properties of our CoarsenRank are summarized as follows: (1) CoarsenRank is designed for mild model misspecification, which assumes there exist the ideal preferences (consistent with model assumption) that locate in a neighborhood of the actual preferences. (2) CoarsenRank then performs regular RAs over a neighborhood of the preferences instead of the original data set directly. Therefore, CoarsenRank enjoys robustness against model misspecification within a neighborhood. (3) The neighborhood of the data set is defined via their empirical data distributions. Further, we put an exponential prior on the unknown size of the neighborhood, and derive a much-simplified posterior formula for CoarsenRank under particular divergence measures. (4) CoarsenRank is further instantiated to Coarsened Thurstone, Coarsened Bradley-Terry, and Coarsened Plackett-Luce with three popular probability ranking models. Meanwhile, tractable optimization strategies are introduced with regards to each instantiation respectively. In the end, we apply CoarsenRank on four real-world data sets. Experiments show that CoarsenRank is fast and robust, achieving consistent improvements over baseline methods.

Keywords: Robust Rank Aggregation, Model Misspecification, CoarsenRank, Coarsened Bradley-Terry, Coarsened Plackett-Luce

1. Introduction

Rank aggregation (RA) refers to the task of recovering the total order over a set of items, given a collection of pairwise/partial/full preferences over items (Lin, 2010). Therefore, RA is a practical and useful approach to summarize user preferences. Preferences could arise not only by explicitly querying users but also through passive data collection, i.e., by observing user purchasing behavior (Baltrunas et al., 2010), clicks on search engine results (Dwork et al., 2001), etc. Compared to rating items, the preferences are more natural expressions of user opinions which can provide more consistent results (Raman and Joachims, 2014). The flexible collection of preferences enables successful application of rank aggregation in various fields, from image rating (Liang and Grauman, 2014) to document recommendation (Sellamanickam et al., 2011), peer grading (Raman and Joachims, 2014), opinion analysis (Chatterjee et al., 2018), bioinformatics (Kim et al., 2014) and mental fatigue monitoring (Pan et al., 2020, 2021).

A basic assumption underlying the vanilla RA is that all preferences are provided by homogeneous users, sharing the same annotation accuracy and agreeing with the single ground truth ranking (Dwork et al., 2001; Li et al., 2017; Chiang et al., 2017; Li et al., 2019). However, this homogeneity assumption is rarely satisfied due to the flexible data construction and the complex real-world situation (Gormley and Murphy, 2005; Kolde et al., 2012; Mollica and Tardella, 2017; Li et al., 2018). For example, the reliability of each user may not be necessarily the same due to the diverse background of each user regarding the candidate items. Moreover, the users are more likely coming from a heterogeneous community and the single total order assumption is no longer suitable since different users judge the items from different perspectives. Therefore, RA usually suffers from model misspecification, namely the inconsistency between the collected ranking data and the homogeneity assumption of RA (Pan et al., 2018).

To address the above inconsistency issue, existing robust RAs resort to an augmentation of the ranking model to account for additional perturbation, where the collected preferences are viewed as a noisy perturbation of some idealized preferences (Han et al., 2018). Particularly, Chen et al. (2013) studied RA in a crowdsourcing environment and proposed CrowdBT for noisy pairwise preferences. CrowdBT models the user reliability with one extra parameter, following the two-coin Dawid-Skene

model (Raykar et al., 2010). Han et al. (2018) proposed ROPAL, which extended CrowdBT for noisy partial preference using more parameters. Note that ROPAL requires a few ground truth preferences from each user for initializing the parameters. It constraints their method to a crowdsourcing setting, where multiple preferences from each user are available. Raman and Joachims (2014) introduced a general framework, called PeerGrader, to aggregate ordinal peer gradings from peer graders while exploring each grader’s reliability by introducing a scale factor. However, each user usually provides one preference in real applications, which would cause overfitting since it needs to estimate the reliability w.r.t. each preference (Sajjadi et al., 2016). The same problem also arises in Xu et al. (2017), which formulates the robust RA as outlier detection and introduces a deviation factor for each preference to account for the unknown noise-perturbation. Indeed, these previous attempts simply amount to convolving the original ranking model with some pre-assumed perturbation mechanism. It leads to a new model with a few more parameters but is just as bound to be misspecified w.r.t. other overlooked perturbations.

The above analysis motivates us to present a novel robust RA approach, called CoarsenRank. The main idea of CoarsenRank is to perform regular RA over a neighborhood of the collected preferences, which enables CoarsenRank against mild model misspecification within the defined neighborhood (Volpi et al., 2018; Chen and Paschalidis, 2018b). However, it is usually intractable to infer directly over the neighborhood of the ranking data because of the unlimited samples involved. Further, it also prohibits sampling-based stochastic gradient solutions in the optimization community due to the particularity of the ranking data. Inspired by Miller and Dunson (2019), which avoids inferring over the neighborhood of the data set by transforming the problem into a tractable fractional likelihood formulation (Bhattacharya et al., 2019). For the sake of tractability, the neighborhood of the data set is first defined as the neighborhood of its empirical data distribution. In particular, the relative entropy is adopted as the divergence metric due to its simplicity (Ben-Tal et al., 2013; Namkoong and Duchi, 2017). We further introduce a prior distribution for the unknown size of the neighborhood to avoid parameter tuning and derive a much-simplified formula for CoarsenRank.

More precisely, we summarize our main contributions in the following:

- We introduce a novel robust rank aggregation method called CoarsenRank. CoarsenRank performs RA over the neighborhood of the ranking data instead of original data set directly. To our best knowledge, CoarsenRank is the first rank aggregation method against model misspecification and enjoys distributional robustness.
- We obtain a computationally efficient formula for CoarsenRank, which introduces only one extra hyperparameter to vanilla ranking models. Further, we instantiate CoarsenRank with three popular probability ranking models and analyze the optimization strategies, respectively. To avoid hyperparameter tuning, an efficient model selection method is introduced to choose the single hyperparameter in a data-driven manner.
- We successfully applied our CoarsenRank on four real-world data sets. Empirical results demonstrate (1) CoarsenRank shows superior reliability against agnostic noises over existing robust RA methods, especially when the number of annotations per user is insufficient; (2) CoarsenRank enjoys linear algorithm complexity, which has great potential for a large-scale scenario.

The rest of this paper is organized as follows. Section 2 discusses RA under model misspecification and introduces the main idea of our CoarsenRank. In Section 3, we pave the theoretical foundation

Notation	Explanation
M	number of items
N	number of preferences
k	length of preferences, which could be variant with regards to each preference
\mathcal{O}	set of items, $\mathcal{O} = \{o_1, o_2, \dots, o_M\}$
$o_i > o_j$	item o_i is preferred over item o_j
ρ_n	n -th real preference, $\rho_n : \rho_n^1 > \rho_n^2 > \dots > \rho_n^k, \{\rho_n^1, \rho_n^2, \dots, \rho_n^k\} \subseteq \mathcal{O}$
ϱ_n	n -th idealized preference, $\varrho_n : \varrho_n^1 > \varrho_n^2 > \dots > \varrho_n^k, \{\varrho_n^1, \varrho_n^2, \dots, \varrho_n^k\} \subseteq \mathcal{O}$
\mathcal{R}_N	collection of real preferences, $\mathcal{R}_N = \{\rho_1, \rho_2, \dots, \rho_N\}$
\mathfrak{R}_N	collection of idealized preferences, $\mathfrak{R}_N = \{\varrho_1, \varrho_2, \dots, \varrho_N\}$, satisfying the homogeneity assumption
$B_s(\mathcal{R}_N, \epsilon)$	sample-level neighborhood of the ranking data set \mathcal{R}_N
$B_d(\mathcal{R}_N, \epsilon)$	distribution-level neighborhood of the ranking data set \mathcal{R}_N
P_θ	probability ranking model, θ denotes the model parameter
P_o	real preference generation distribution
$F_N(x \mathcal{R}_N)$	the empirical distributions of the real ranking data set \mathcal{R}_N
$F_N(x \mathfrak{R}_N)$	the empirical distributions of the idealized ranking data set \mathfrak{R}_N
Θ	space of possible parameter values that defines a ranking model P_θ
\mathcal{P}	set of all probability rank models with parameterization space Θ , $P_\theta \in \mathcal{P}$ and $\theta \in \Theta$
$D(\cdot, \cdot)$	divergence metric between two data sets, defined via their empirical data distributions
$\mathbb{I}_x(y)$	indicator function, which is one at $x = y$ or zero otherwise

Table 1: Common mathematical notations

for CoarsenRank and illustrate how CoarsenRank enables us to perform robust RA against model misspecification. Section 4 presents an efficient EM algorithm as well as a Gibbs sampling algorithm for CoarsenRank and discusses a data-driven strategy for hyperparameter selection. Section 5 summarizes the differences between our CoarsenRank and related (robust) RA models. Section 6 demonstrates the efficacy of CoarsenRank through empirical results on four real-world data sets. Section 7 concludes the paper and envisions future work.

2. Problem Statement and Literature Review

In this section, we first introduce the problem setting of vanilla RA and RA under model misspecification. Furthermore, we summarize previous robust RA for alleviating the model misspecification, as well as their deficiencies. Then, we motivate our Coarsened RA, which perform regular RA over a neighborhood of the collected preferences, and therefore enjoys distributional robustness against noise-agnostic perturbation within a neighborhood.

2.1 Rank Aggregation

In Table 1, we first illustrate the common mathematical notations that are used later.

Let \mathcal{R}_N denote a collection of partial preferences $\{\rho_1, \rho_2, \dots, \rho_N\}$ over the item set $\mathcal{O} = \{o_1, o_2, \dots, o_M\}$. The goal of rank aggregation is then to aggregate the collected preferences \mathcal{R}_N

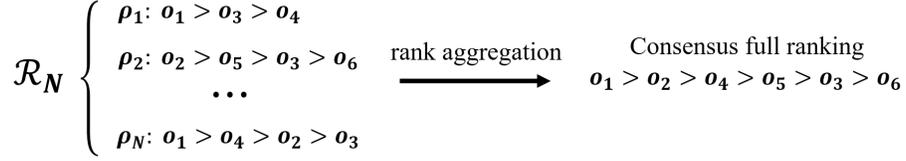


Figure 1: The paradigm of rank aggregation.

into a consensus order over all M items in O (See Figure 1). The consensus order should achieve the maximum agreement among all preferences in \mathcal{R}_N (Dwork et al., 2001).

In this paper, we focus our work on rank aggregation using the probability ranking model. Particularly, it assumes there exists a preference generative model P_o from which the preferences \mathcal{R}_N are sampled, i.e., $\mathcal{R}_N = \{\rho_n | \rho_n \sim P_o, n = 1, 2, \dots, N\}$. However, the real data generation model P_o is hardly accessible due to the complexity of the real situation. For the sake of easier modeling, a parameterized rank model P_θ is usually adopted under the assumption of homogeneity of users¹. Let \mathcal{P} be the set of all probability rank models under the homogeneity assumption. For the sake of easier modeling, a parameterized rank model P_θ is usually adopted under the assumption of homogeneity of users. Then a maximum likelihood estimation (MLE) for RA can be formulated as follows,

$$\max_{\theta \in \Theta} P_\theta(\mathcal{R}_N), \quad \text{where } P_\theta \in \mathcal{P} \text{ and } \mathcal{R}_N = \{\rho_n | \rho_n \sim P_o, n = 1, 2, \dots, N\}. \quad (1)$$

$P_\theta(\mathcal{R}_N) = \prod_{n=1}^N P_\theta(\rho_n)$ denotes the likelihood over the collected preferences \mathcal{R}_N . P_θ is usually instantiated with Thurstone model (Thurstone, 1927a,b), Bradley-Terry model (Bradley and Terry, 1952), Plackett-Luce model (Plackett, 1975; Luce, 1959), etc. Note that the model parameter θ is usually associated with each item, where the full ranking list could be derived accordingly after θ is inferred. For example, the full ranking list can be obtained by sorting the model parameter θ in the case of the Thurstone/Bradley-Terry/Plackett-Luce model.

2.2 Rank Aggregation under Model Misspecification

In this section, we discuss RA under model misspecification. The term ‘‘model misspecification’’ here refers to the mismatch between the ranking model P_θ and the ranking data set \mathcal{R}_N , namely the collected user preferences do not strictly satisfy the user homogeneity assumption of the ranking model.

The model misspecification would arise when preferences were not strictly collected from a homogeneous user community due to the flexible data construction and the complex real situation (See Figure 2). For example, the reliability of each user would not be the same and the single total order assumption would be no longer satisfied. Mathematically, we adopt the parameterized ranking model $P_\theta \in \mathcal{P}$ under the homogeneity assumption, while the real preference generation distribution P_o violates this assumption, i.e., $P_o \notin \mathcal{P}$. Therefore, an MLE for RA under model misspecification can be formulated as follows,

$$\max_{\theta \in \Theta} P_\theta(\mathcal{R}_N) \quad \text{where } P_\theta \in \mathcal{P}, P_o \notin \mathcal{P} \text{ and } \mathcal{R}_N = \{\rho_n | \rho_n \sim P_o, n = 1, 2, \dots, N\}. \quad (2)$$

1. Sampling the partial preferences from a specific probability ranking model is not our focus in this paper. Please refer to Liu et al. (2019); Zhao and Xia (2019) for related literature.

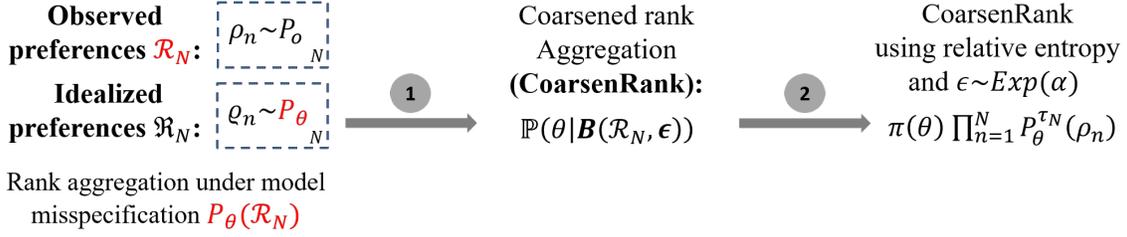


Figure 2: The logic stream of our CoarsenRank. Stage 1: performing rank aggregation over a neighborhood of the collected preferences (See Equation (7) and Equation (8)). Stage 2: adopting relative entropy as the divergence measure and assign an exponential prior for the size of the neighborhood (See Theorem 2 and Theorem 3).

For the sake of explanation, let \mathfrak{R}_N represent a virtual data set $\{\varrho_n | \varrho_n \sim P_\theta, n = 1, 2, \dots, N\}$, which consists of idealized preferences and satisfies the homogeneity assumption. Then, RA under model misspecification can be formulated as noisy RA, where the collected preferences are viewed as a noisy perturbation of some idealized preferences.

Then, we come to robust rank aggregation against model misspecification, namely how to achieve a reliable total order from the collected preferences \mathcal{R}_N using a misspecified ranking model P_θ .

2.3 Previous Attempts: Convolving Ranking Model with Certain Perturbation Mechanisms

When encountering model misspecification, a remedy solution for accessing a correct rank model could be

$$P_\theta(\rho_n) = \sum_{\varrho_n} P_\theta(\varrho_n) P(\rho_n | \varrho_n) = \sum_{\varrho_n} P(\rho_n, \varrho_n), \quad \text{where } \rho_n \sim P_o, \varrho_n \sim P_\theta. \quad (3)$$

Considering the discrete characteristics of ranking space, the ranking distribution i.e., P_o and P_θ as well as the conditional distribution $P(\rho_n | \varrho_n)$ and joint distribution $P(\rho_n, \varrho_n)$ should all be the discrete distribution.

Previous approaches usually resort to an augmentation of the ranking model to account for additional error/noise/uncertainty caused misspecification. For the sake of tractability, the perturbation mechanism is usually defined at the sample level. According to Equation (3), there are essentially two ways of implementing this:

- One intuitive approach is to correct each preference by pre-assuming some perturbation distribution, i.e., $P(\rho_n | \varrho_n)$. However, this simply amounts to convolving the original model distribution P_θ with the predefined perturbation, leading to a new model that has a few more parameters but is just as bound to be misspecified w.r.t. other overlooked perturbations.
- The second approach would be to model the joint distribution $P(\rho_n, \varrho_n)$ directly, which needs to take into consideration all potential perturbations. Essentially, it needs to be a nonparametric model for $P(\rho_n, \varrho_n)$, but would easily be computationally intractable.

Meanwhile, the perturbation patterns leading to model misspecification vary from setting to setting. It is impossible to design a universal practice that can be generalized to most settings. Therefore, in this paper, we perform rank aggregation against model misspecification from another perspective.

2.4 Our CoarsenRank: Rank Aggregation over The Neighborhood of Ranking Data

In many situations, it is impractical to correct the model, and these are the situations our method is intended to address. We are concerned with robust rank aggregation against model misspecification (Equation (2)) in general, not just one particular kind of perturbation considered in previous work. Recent advances of robust Bayesian inference (Miller and Dunson, 2019; Volpi et al., 2018) raised the Coarsening mechanism, namely *inferring over the neighborhood of the original data set would equip the learning model with distributional robustness*. Motivated by the proposed Coarsening mechanism, we consider performing rank aggregation over the neighborhood of the ranking data.

To deliver our model, we first give the definition of the neighborhood in the sense of ranking data as follows,

Definition 1 (sample-level neighborhood) Let \mathcal{R}_N denote the ranking data set and ρ represent one preference belonging to \mathcal{R}_N . We define the neighborhood $B_s(\mathcal{R}_N, \epsilon)$ of the \mathcal{R}_N with size $\epsilon > 0$ as follows:

$$B_s(\mathcal{R}_N, \epsilon) = \{\rho' | D(\rho', \rho) < \epsilon, \exists \rho \in \mathcal{R}_N\}, \quad (4)$$

where $D(\cdot, \cdot)$ denotes some distance measure between two preferences ρ' and ρ , e.g., Kendall tau distance (Kendall, 1938), Spearman's rank correlation (Ramsey, 1989).

Definition 2 (distribution-level neighborhood) Let \mathcal{R}_N denote the ranking data set and $B_d(\mathcal{R}_N, \epsilon)$ denote the neighborhood of the ranking data set \mathcal{R}_N with size $\epsilon > 0$. Then, we define

$$B_d(\mathcal{R}_N, \epsilon) = \{\mathcal{R}'_N | D(\mathcal{R}'_N, \mathcal{R}_N) < \epsilon\}, \quad (5)$$

where $D(\cdot, \cdot)$ denotes some distance measures between two ranking data sets. The distance measure between two data sets is usually defined as the divergence of their corresponding empirical distributions. Popular divergence measures between distributions are Kullback-Leibler (KL) divergence (Kullback and Leibler, 1951), f -divergence (Ali and Silvey, 1966) and Wasserstein metric (Villani, 2008).

Proposition 1 Given any ranking data set \mathcal{R}_N , \mathcal{R}_N must be (1) a subset of its sample-level neighborhood $B_s(\mathcal{R}_N, \epsilon)$ if $D(\cdot, \cdot)$ is defined between two preferences, or (2) an element of its distribution-level neighborhood $B_d(\mathcal{R}_N, \epsilon)$ if $D(\cdot, \cdot)$ is defined between two ranking data sets. Namely

$$\begin{aligned} \mathcal{R}_N &\subseteq B_s(\mathcal{R}_N, \epsilon), & \text{sample-level neighborhood,} \\ \mathcal{R}_N &\in B_d(\mathcal{R}_N, \epsilon), & \text{distribution-level neighborhood.} \end{aligned}$$

Proof: In terms of sample-level neighborhood, we have

$$\forall \rho \in \mathcal{R}_N, \quad \exists \rho^* = \rho, \text{ s.t. } D(\rho^*, \rho) = 0 < \epsilon.$$

Then $\rho \in B_s(\mathcal{R}_N, \epsilon)$ holds. Accordingly, $\mathcal{R}_N \subseteq B_s(\mathcal{R}_N, \epsilon)$ holds according to the definition of sample-level neighborhood in Equation (4).

In terms of distribution-level neighborhood, we have

$$D(\mathcal{R}_N, \mathcal{R}_N) = 0 < \epsilon.$$

Therefore, we have $\mathcal{R}_N \in B_d(\mathcal{R}_N, \epsilon)$ hold according to the definition of distribution-level neighborhood in Equation (5).

Note that the proof is valid for any particular choice of the distance metric, either sample level or distribution level. ■

Definition 3 (empirical data distribution) Let $F_N(x|\mathcal{R}_N)$ and $F_N(x|\mathfrak{R}_N)$ denote the empirical distributions of the ranking data sets \mathcal{R}_N and \mathfrak{R}_N , respectively.

$$\begin{aligned} F_N(x|\mathcal{R}_N) &= \frac{1}{N} \sum_{n=1}^N \mathbb{I}_{\rho_n}(x), \quad \text{where } \mathbb{I}_{\rho_n}(x) = \begin{cases} 1, & x = \rho_n \\ 0, & x \neq \rho_n \end{cases}, \\ F_N(x|\mathfrak{R}_N) &= \frac{1}{N} \sum_{n=1}^N \mathbb{I}_{\varrho_n}(x), \quad \text{where } \mathbb{I}_{\varrho_n}(x) = \begin{cases} 1, & x = \varrho_n \\ 0, & x \neq \varrho_n \end{cases}. \end{aligned} \quad (6)$$

In this paper, we assume that the empirical distribution converges to the corresponding preference generation distribution, namely $F_N(x|\mathcal{R}_N) \rightarrow P_o$ and $F_N(x|\mathfrak{R}_N) \rightarrow P_\theta$ when $N \rightarrow +\infty$.

For the sake of brevity, we introduce our work following the definition of distribution-level neighborhood. Let $\mathfrak{R}_N \sim P_\theta$ denote every preference of the ranking data set \mathfrak{R}_N is sampled from the ranking distribution P_θ , namely $\forall \varrho_n \in \mathfrak{R}_N, \varrho_n \sim P_\theta$. We assume that the idealized ranking data set \mathfrak{R}_N locates in the small neighborhood of the actually collected preferences \mathcal{R}_N , i.e.,

$$\exists \text{ a small } \epsilon, \forall \mathfrak{R}_N \sim P_\theta, \text{ we have } \mathfrak{R}_N \in B_d(\mathcal{R}_N, \epsilon)$$

This is a basic assumption in distributional robustness literature (Chen and Paschalidis, 2018b). Otherwise, if the size of neighborhood ϵ which satisfies our assumption is very large, it means a completely wrong model is adopted and it is impossible to learn a meaningful result. This is why we call our setting ‘‘mild model misspecification’’. Meanwhile, the sense of ‘‘neighborhood’’ in the distribution level covers most types of noise perturbations (Chen and Paschalidis, 2018a). Therefore, the MLE of our Coarsened rank aggregation (CoarsenRank) can be formulated as follows,

$$\max_{\theta \in \Theta} P_\theta(\mathfrak{R}_N), \quad \text{where } \mathfrak{R}_N \sim P_\theta \text{ and } \mathfrak{R}_N \in B_d(\mathcal{R}_N, \epsilon). \quad (7)$$

Note that we use the word ‘‘Coarsen’’ to emphasize the learning paradigm which pursues the distributional robustness by inferring over the neighborhood of the data set (Miller and Dunson, 2019).

An equivalent (but compact) formulation of our CoarsenRank (Equation (7)) can be derived as follows

$$\begin{aligned} \text{Equation (7)} &\stackrel{\textcircled{1}}{\iff} \max_{\theta \in \Theta} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} \mathbb{P}(\mathfrak{R}_N | \theta, D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon) \\ &\stackrel{\textcircled{2}}{\implies} \max_{\theta \in \Theta} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} \mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon), \end{aligned} \quad (8)$$

where $\textcircled{1}$ is an equivalent MLE formulation rearranged according to the definition of distribution-level neighborhood (Equation (5)). By assigning a suitable prior for model parameter θ , i.e., $\theta \sim \pi(\theta)$, $\textcircled{2}$ deduces the maximum a posteriori probability (MAP) estimate of our CoarsenRank model.

Equation (8) reveals that: (1) our CoarsenRank degenerates to vanilla rank aggregation method (Equation (1)) when the collected preferences satisfy the homogeneity assumption; (2) our CoarsenRank would be robust to noise-agnostic perturbations as long as the idealized preferences locates in the neighborhood of the collected preferences when model misspecification arises; and (3) our CoarsenRank would fail to output a reliable total ranking list when the collected preferences

significantly violate the homogeneity assumption. The same is true for other vanilla rank aggregation methods and most of robust RAs which fail to capture this perturbation. Therefore, compared with the previous methods, our CoarsenRank is robust to most potential perturbations within a neighborhood, not only to some pre-assumed perturbations.

Remark 1 (Coarsening Mechanism VS. Minimax Distributional Robustness) *The Coarsening mechanism $\max_{\theta \in \Theta} \mathbb{E}_{\mathfrak{R}_N \in B_d(\mathcal{R}_N, \epsilon)} P_\theta(\mathfrak{R}_N)$ shares a similar formula with the minimax distributional robustness $\min_{\theta \in \Theta} \max_{\mathfrak{R}_N \in B_d(\mathcal{R}_N, \epsilon)} \ell(\mathfrak{R}_N | \theta)$ (Sinha et al., 2017). $\ell(\theta)$ denotes the loss function, which is usually replaced with the negative log-likelihood. The Coarsening mechanism aims to maximize the likelihood over the neighborhood of original data set \mathcal{R}_N . The minimax distributional robustness aims to minimize the loss using the worst data samples in the neighborhood of original data set \mathcal{R}_N . The neighborhood in the Coarsening mechanism is usually defined at the distribution-level (Equation (5)) where a Bayesian criterion can be adopted to estimate the proper size of the neighborhood for each data set; while the neighborhood in the minimax distributional robustness is usually defined at the sample-level (Equation (5)) and a fixed size neighborhood is adopted once for all. The choice of distance measures $D(\cdot, \cdot)$ influences robustness guarantee and tractability in the both two paradigms. ■*

Remark 2 (Coarsening Mechanism VS. Rank-dependent Coarsening) *The word Coarsening also rises in Fahandar et al. (2017), which, however, has a totally different meaning there. The term “rank-dependent coarsening” refers to the process of turning a full ranking into an incomplete one (Fahandar et al., 2017). It is different from our “Coarsening mechanism”, which refers to the paradigm that performing the Bayesian inference over the neighborhood of the original data set. ■*

3. Coarsened Rank Aggregation

In this section, we first illustrate how CoarsenRank enables us to perform robust rank aggregation against model misspecification. Meanwhile, a simplified formula is derived for CoarsenRank, which introduces only one extra hyperparameter to vanilla ranking models. Then, we instantiate our CoarsenRank framework with three popular probability ranking models and analyze their optimization strategies, respectively.

3.1 Distributional Robustness of Coarsening Mechanism

Assuming that the empirical distribution defined in Equation (6) converges to the corresponding data generating distribution, namely $F_N(x|\mathcal{R}_N) \rightarrow P_o$ and $F_N(x|\mathfrak{R}_N) \rightarrow P_\theta$ when $N \rightarrow +\infty$, we come to Theorem 1. Note that this result is essentially S3.1 in the Supplement of Miller and Dunson (2019). We also provide a proof in the Appendix for the sake of completeness.

Theorem 1 *Suppose $D(\mathcal{R}_N, \mathfrak{R}_N)$ is an almost surely-consistent estimator² of $D(P_o, P_\theta)$, namely $D(\mathcal{R}_N, \mathfrak{R}_N) \xrightarrow[N \rightarrow +\infty]{a.s.} D(P_o, P_\theta)$, where $F_N(x|\mathcal{R}_N) \rightarrow P_o$ and $F_N(x|\mathfrak{R}_N) \rightarrow P_\theta$ when $N \rightarrow +\infty$. Assume $\mathbb{P}(D(P_o, P_\theta) = \epsilon) = 0$ and $\mathbb{P}(D(P_o, P_\theta) < \epsilon) > 0$, then we have*

$$\mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon) \xrightarrow[N \rightarrow +\infty]{a.s.} \mathbb{P}(\theta | D(P_o, P_\theta) < \epsilon), \quad (9)$$

for any $\theta \in \Theta$ such that $\int |\theta| \mathbb{P}(d\theta) < \infty$.

2. In probability theory, an event happens almost surely if it happens with probability one.

Theorem 1 is a general conclusion in robust Bayesian inference (Miller and Dunson, 2019). It justifies our motivation to pursue robustness in a distributional sense. In what follows, we extend Theorem 1 to some variants which possess nice properties for robust rank aggregation.

3.1.1 LEVEL OF DISTRIBUTIONAL ROBUSTNESS

The value of the parameter ϵ denotes the level of deviation about the actually collected preferences from the idealized preferences, which varies from data set to data set. Simply fixing the ϵ to a small value, the Coarsening mechanism degenerates to a minimum-expectation problem as we discussed in Remark 1. It would be heuristic without sufficient prior knowledge about idealized preferences, since a small ϵ may fail to account for the unknown distribution deviation while a large ϵ means an exponential level of ranking space to search. To ease the burden of pre-defining ϵ , we treat it as a random variable and introduce a prior on it, where an efficient model selection method is introduced. In particular, we have the following conclusion.

Theorem 2 *Assume $\theta \sim \pi(\theta)$, the approximate posterior can be further simplified.*

$$\mathbb{P}(\theta | D(P_o, P_\theta) < \epsilon) = \frac{\pi(\theta) \mathbb{P}(D(P_o, P_\theta) < \epsilon | \theta)}{\int_{\theta} \pi(\theta) \mathbb{P}(D(P_o, P_\theta) < \epsilon | \theta) d\theta} \propto \exp(-\alpha D(P_o, P_\theta)) \pi(\theta), \quad (10)$$

when random variable ϵ subjects to an exponential prior, i.e., $\epsilon \sim \text{Exp}(\alpha)$.

Proof: Note that since $\epsilon \sim \text{Exp}(\alpha)$, we have

$$\begin{aligned} \mathbb{P}(D(P_o, P_\theta) < \epsilon | \theta) &= 1 - \mathbb{P}(\epsilon \leq D(P_o, P_\theta) | \theta) \\ &= 1 - (1 - \exp(-\alpha D(P_o, P_\theta))) \\ &= \exp(-\alpha D(P_o, P_\theta)), \end{aligned}$$

where the second equation holds because the cumulative distribution function $\mathbb{P}(D(P_o, P_\theta) > \epsilon | \theta)$ is independent of θ (Bishop, 2006). Then, we can substitute $\mathbb{P}(D(P_o, P_\theta) < \epsilon | \theta)$ in Equation (10) with $\exp(-\alpha D(P_o, P_\theta))$ and complete the proof while omitting the normalization constant. ■

Indeed, a very large class of distributions can be adopted as the prior for ϵ . A case of particular interest arises when $\epsilon \sim \text{Exp}(\alpha)$, since it leads to a computationally simple formula via maintaining an exponential formulation. The efficacy of the exponential prior is verified in our experiment (See Section 6).

Inspired by the exponential formulation of the posterior derived in Equation (10), we give the following derivations (Equation (11)) to explain why the vanilla rank aggregation is lack of robustness.

$$\begin{aligned} \mathbb{E}_{\mathfrak{R}_N \in B_d(\mathcal{R}_N, 0)} \mathbb{P}(\theta | \mathfrak{R}_N) &= \mathbb{P}(\theta | \mathcal{R}_N) = \frac{\pi(\theta) P_\theta(\mathcal{R}_N)}{\int \pi(\theta) P_\theta(\mathcal{R}_N) d\theta} \propto \pi(\theta) P_\theta(\mathcal{R}_N) \quad (11) \\ &\stackrel{\textcircled{1}}{=} \pi(\theta) \exp\left(\sum_{n=1}^N \log P_\theta(\rho_n)\right) \stackrel{\textcircled{2}}{=} \pi(\theta) \exp\left(N \sum_{n=1}^N F_N(\rho_n | \mathcal{R}_N) \log P_\theta(\rho_n)\right) \\ &\stackrel{\textcircled{3}}{\approx} \pi(\theta) \exp\left(N \sum P_o \log P_\theta\right) \stackrel{\textcircled{4}}{\propto} \pi(\theta) \exp(-N \mathcal{D}_{\text{KL}}(P_o \| P_\theta)), \end{aligned}$$

where $\textcircled{1}$ holds because $P_\theta(\mathcal{R}_N) = \prod_{n=1}^N P_\theta(\rho_n)$ is the likelihood. $\textcircled{2}$ holds following the definition of the empirical data distribution $F_N(x | \mathcal{R}_N) = \frac{1}{N} \sum_{n=1}^N \mathbb{I}_{\rho_n}(x)$. $\textcircled{3}$ indicates Monte Carlo

approximation. ④ holds due to the added entropy term $\sum P_o \log P_o$, which is a constant w.r.t. the model parameter θ . The standard posterior (Equation (11)) tends to zero under model misspecification ($P_o \neq P_\theta$) as $N \rightarrow +\infty$, while the approximate posterior (Equation (10)) remains stable. This verifies our motivation for pursuing robust RA since vanilla RA, as well as data-augmentation based RAs, would inevitably output unreliably results even when infinity samples are available.

3.1.2 TYPES OF DISTRIBUTIONAL ROBUSTNESS AND TRACTABILITY

The choice of $D(\cdot, \cdot)$ in $\mathbb{P}(\theta|D(P_o, P_\theta) < \epsilon)$ (Equation (9)) affects both the richness of the robustness types as well as the tractability of the resultant optimization problem. The Wasserstein metric is a popular option in previous approaches on distributional robustness (Blanchet et al., 2016; Gao et al., 2017; Volpi et al., 2018), which exhibits superior tolerance to adversarially corrupted outliers (Chen and Paschalidis, 2018a,b) and also allows robustness to unseen data (Abadeh et al., 2015; Sinha et al., 2017). Meanwhile, Ben-Tal et al. (2013); Namkoong and Duchi (2017) adopted f -divergences in pursuit of tractable optimization approaches. It is worthy noting that the rank aggregation task has its particularities. First, no generalization test is required for the RA task since we only need to aggregate the whole ranking data set into one consensus full rank. Second, the probability ranking model itself has high complexity. Therefore, we consider relative entropy for $D(\cdot, \cdot)$, since it allows standard inference with no additional computational burden and helps to exhibit robustness to most types of perturbations.

Before introducing Theorem 3, we first introduce Lemma 1 which contains some preliminary results from Miller and Dunson (2019).

Lemma 1 (Miller and Dunson (2019)) *Let $\Delta_d = \{p \in \mathbb{R}^d : \sum_i^d p_i = 1, p_i > 0 \forall i\}$, and $q \in \Delta_d$. We argue that if x_1, x_2, \dots, x_N i.i.d. $\sim q$ and $F_N(t|x_{1:N}) = \frac{1}{N} \sum_{n=1}^N \mathbb{I}_{x_n}(t)$, then for $p \in \Delta_k$ near q in KL divergence,*

$$\mathbb{E}_{x_{1:N} \sim q} [\exp(-\alpha \mathcal{D}_{\text{KL}}(p \| F_N(t|x_{1:N})))] \approx \left(\frac{N\tau_N}{\alpha} \right)^{\frac{k-1}{2}} \exp(-N\tau_N \mathcal{D}_{\text{KL}}(p \| q)),$$

where $\tau_N = \frac{1/N}{1/N+1/\alpha}$. ■

Lemma 1 is defined for discrete distribution, which can be applied to our probability ranking model. To be specific, in terms of the item set $\mathcal{O} = \{o_1, o_2, \dots, o_M\}$, there are totally $\sum_{i=2}^M \binom{M}{i} i!$ possible ranking lists. Therefore, the probability ranking model P_θ is actually a discrete distribution with $\sum_{i=2}^M \binom{M}{i} i!$ supports.

Theorem 3 *Suppose relative entropy is adopted as the distance measure, namely $D(\mathcal{R}_N, \mathfrak{R}_N) = \mathcal{D}_{\text{KL}}(F_N(x|\mathcal{R}_N) \| F_N(x|\mathfrak{R}_N)) = \int F_N(x|\mathcal{R}_N) \log \frac{F_N(x|\mathcal{R}_N)}{F_N(x|\mathfrak{R}_N)}$, and the empirical distribution converges to the corresponding data generating distribution, namely $F_N(x|\mathcal{R}_N) \rightarrow P_o$ and $F_N(x|\mathfrak{R}_N) \rightarrow P_\theta$ when $N \rightarrow +\infty$. If ϵ is subject to an exponential prior, i.e., $\epsilon \sim \text{Exp}(\alpha)$, we can obtain the following simple approximation to our CoarsenRank in Equation (8):*

$$\max_{\theta \in \Theta} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon)] \propto \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N P_\theta^{\tau_N}(\rho_n), \quad (12)$$

where \propto denotes that the term on the left is approximately equal to a term, which is proportional to the expression on the right, and $\tau_N = \frac{1/N}{1/N+1/\alpha}$.

Proof: According to Theorem 2, we have

$$\begin{aligned} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon)] &\propto \pi(\theta) \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\mathbb{P}(D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon) | \theta] \\ &= \pi(\theta) \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\exp(-\alpha D(\mathcal{R}_N, \mathfrak{R}_N))], \end{aligned}$$

where we omit the normalization constant with respect to θ .

Further, we have

$$\begin{aligned} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\exp(-\alpha D(\mathcal{R}_N, \mathfrak{R}_N))] &\stackrel{\textcircled{1}}{=} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\exp(-\alpha \mathcal{D}_{\text{KL}}(F_N(x|\mathcal{R}_N) \| F_N(x|\mathfrak{R}_N)))] \\ &\stackrel{\textcircled{2}}{\asymp} \exp(-N\tau_N \mathcal{D}_{\text{KL}}(F_N(x|\mathcal{R}_N) \| P_\theta)) \\ &\stackrel{\textcircled{3}}{\propto} \exp\left(N\tau_N \sum_{n=1}^N F_N(\rho_n|\mathcal{R}_N) \log P_\theta(\rho_n)\right) \stackrel{\textcircled{4}}{=} \prod_{n=1}^N P_\theta^{\tau_N}(\rho_n), \end{aligned}$$

where $\tau_N = \frac{1/N}{1/N+1/\alpha}$. $\textcircled{1}$ is valid by instantiating the distance measure with relative entropy. $\textcircled{2}$ follows Lemma 1 while omitting the constant-coefficient. $\textcircled{3}$ holds due to the removal of the constant entropy term $\sum_{n=1}^N F_N(\rho_n|\mathcal{R}_N) \log F_N(\rho_n|\mathcal{R}_N)$, which is a constant w.r.t. the model parameter θ . $\textcircled{4}$ holds according to the definition of the empirical data distribution $F_N(x|\mathcal{R}_N) = \frac{1}{N} \sum_{n=1}^N \mathbb{I}_{\rho_n}(x)$. ■

Remark 3 (Connection between CoarsenRank and the standard posterior) Since $\epsilon \sim \text{Exp}(\alpha)$, we have $\mathbb{E}(\epsilon) = \frac{1}{\alpha}$ denoting the expected discrepancy of the collected preferences \mathcal{R}_N w.r.t. \mathfrak{R}_N . Further, $\mathbb{E}(\epsilon)$ tends to zero as $\alpha \rightarrow +\infty$, which means the misspecification does not exist in the limit. Accordingly, the robust posterior Equation (12) degenerates to the standard posterior as $\tau_N = \frac{1/N}{1/N+1/\alpha}$ approximates to 1 when $\alpha \rightarrow +\infty$. ■

3.2 Instantiating CoarsenRank with Various Probability Ranking Model

Based on the CoarsenRank framework (Equation (12)), we instantiate P_θ with various probability ranking models.

Please note an actual probability ranking model is usually defined over a subset of items and satisfies the summing up to 1 condition only on this subset. However, it does not necessarily satisfy the definition of the probability distribution over the whole item set \mathcal{O} , namely the probability integration of a probability ranking model over \mathcal{O} is not 1 (Zhao and Xia, 2019)³. Without loss of generality, we assume each subset of items is uniformly sampled from the whole set. Then, the normalization constant can be safely omitted during optimization without affecting our final estimation.

3.2.1 COARSENEDED THURSTONE MODEL (COARSENTH)

We first review the basic Thurstone model (Thurstone, 1927b), which is a popular ranking model to model pairwise comparisons. Particularly, it assumes that the score $\theta_m \in \mathbb{R}$ for each item o_m follows a Gaussian distribution $N(\mu_m, \sigma_m^2)$, $\forall m = 1, 2, \dots, M$. For simplicity, we only consider the

3. In terms of pairwise comparisons, there are totally $\binom{M}{2}$ distinct pairs with the corresponding probability integration being $\binom{M}{2}$. In terms of listwise preferences, there are $\sum_{i=2}^M \binom{M}{i}$ distinct subsets in total with the corresponding probability integration being $\sum_{i=2}^M \binom{M}{i}$.

Thurstone model with $\sigma_m = 1$ for all items. In particular, the comparison between any two items o_i and o_j also follows a Gaussian distribution $N(\mu_i - \mu_j, 2)$.

For a pairwise comparison $\rho_n : \rho_n^1 > \rho_n^2$, Thurstone model assumes

$$P_\theta(\rho_n) = \Phi\left(\frac{\Delta\theta_{\rho_n}}{\sqrt{2}}\right) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{\Delta\theta_{\rho_n}}{\sqrt{2}}} \exp\left(-\frac{t^2}{2}\right) dt, \quad (13)$$

where $\Delta\theta_{\rho_n} = \theta_{\rho_n^1} - \theta_{\rho_n^2}$ denotes the difference between the score of the two items in ρ_n . Φ is the cumulative distribution function (CDF) of the standard normal distribution.

According to Theorem 3, an example of our CoarsenRank (Equation (12)) using Thurstone model can be represented as follows:

$$\begin{aligned} \max_{\theta \in \Theta} \mathbb{E}_{\mathfrak{R}_N \sim P_\theta} [\mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{R}_N) < \epsilon)] &\approx \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N P_\theta^{\tau_N}(\rho_n) \\ &= \max_{\theta \in \Theta} \frac{\pi(\theta)}{(\sqrt{2\pi})^{N\tau_N}} \left[\prod_{n=1}^N \int_{-\infty}^{\frac{\Delta\theta_{\rho_n}}{\sqrt{2}}} \exp\left(-\frac{t^2}{2}\right) dt \right]^{\tau_N} \end{aligned} \quad (14)$$

where $\tau_N = \frac{1/N}{1/N+1/\alpha}$. Equation (14) can be only applied to pairwise comparisons. When encountering listwise preferences, we need to split each listwise preferences into pairwise comparisons and then alternatively perform our CoarsenRank on the new data set (Khetan and Oh, 2016).

Remark 4 (Optimization Intractability and Our Strategy) *The cumulative distribution function Φ is a special function, which cannot be expressed in terms of elementary functions. Therefore, it is inefficient or intractable to optimize Equation (14) directly.*

Inspired by the work which explores the connection between the sigmoid function and the cumulative Gaussian distribution (Weng and Lin, 2011), we consider approximating the cumulative distribution function Φ with the sigmoid function. In particular,

$$\Phi\left(\frac{\Delta\theta_{\rho_n}}{\sqrt{2}}\right) \approx \frac{1}{1 + \exp(-\lambda\Delta\theta_{\rho_n})}, \quad (15)$$

where λ is set as $2/\sqrt{\pi}$ so that the two probability curves have the same slope at $\Delta\theta_{\rho_n} = 0$. A more accurate approximation with the second order moments guarantee can be found in Daunizeau (2017). Then, Equation (14) can be further approximated as

$$\text{Equation (14)} \approx \max_{\theta \in \Theta} \frac{\pi(\theta)}{(\sqrt{2\pi})^{N\tau_N}} \prod_{n=1}^N \frac{1}{[1 + \exp(-\lambda\Delta\theta_{\rho_n})]^{\tau_N}}, \quad (16)$$

where regular gradient-based optimization approaches could be carried out. In particular, the score is simply initialized to zero for all items, namely $\theta_m = 0, \forall m = 1, 2, \dots, M$. ■

3.2.2 COARSENEDED BRADLEY-TERRY MODEL (COARSENBT)

A closely related model to the Thurstone model is the Bradley-Terry (BT) model (Bradley and Terry, 1952). For any pairwise comparison $\rho_n : \rho_n^1 > \rho_n^2$, the BT model assumes

$$P_\theta(\rho_n) = \frac{\theta_{\rho_n^1}}{\theta_{\rho_n^1} + \theta_{\rho_n^2}}, \quad (17)$$

where $\theta_m \in R_+^M$ is a positive support parameter for item o_m , $\forall m = 1, 2, \dots, M$.

According to Theorem 3, an example of our CoarsenRank (Equation (12)) using BT model can be represented as follows:

$$\text{Equation (12)} \propto \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N P_{\theta}^{\tau_N}(\rho_n) = \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N \left[\frac{\theta_{\rho_n^1}}{\theta_{\rho_n^1} + \theta_{\rho_n^2}} \right]^{\tau_N}, \quad (18)$$

where $\tau_N = \frac{1/N}{1/N+1/\alpha}$. Similar to Thurstone model, Equation (18) can only model pairwise comparisons. We still adopt the rank breaking strategy to split each listwise preferences into pairwise comparisons and perform our CoarsenRank on the new data set (Khetan and Oh, 2016).

Remark 5 (Optimization Intractability and The Data Augmentation Method) *The main inferential issue related to Equation (18) concerns the presence of the annoying normalization terms $\theta_{\rho_n^1} + \theta_{\rho_n^2}$, $\forall n = 1, 2, \dots, N$, that do not permit the direct maximization of the posterior. Further, the nonnegative constraint over the model parameters θ rules out the direct applications of gradient-based optimization approaches.*

Motivated by Caron and Doucet (2012), we introduce the data augmentation method to address the above-mentioned difficulty. Considering the fact that the Gumbel distribution is employed as a distribution of the support parameters and the conjugacy of the Gamma density with the Gumbel distribution, we follow Caron and Doucet (2012) and introduce an auxiliary Gamma random variable for each normalization term, which leads to a joint distribution without suffering from the annoying normalization terms. ■

3.2.3 COARSENEDED PLACKETT-LUCE MODEL (COARSENPL)

Here we instantiate P_{θ} with the popular Plackett-Luce (PL) model (Plackett, 1975; Luce, 1959). Different from the previous Thurstone mode and the BT model, PL model is a more general probability ranking model, which could model listwise rankings of a finite set of items directly. Note that PL model incorporates BT model as a special case.

For a ranking list $\rho_n : \rho_n^1 > \rho_n^2 > \dots > \rho_n^k$, the PL model assumes

$$P_{\theta}(\rho_n) = \prod_{i=1}^{k-1} \frac{\theta_{\rho_n^i}}{\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}}, \quad (19)$$

where $\theta_m \in R_+^M$ is a positive support parameter associated with item o_m , $\forall m = 1, 2, \dots, M$. Comparing Equation (19) to Equation (17), PL model degenerates to BT model when modeling pairwise comparisons, i.e., $k \equiv 2$.

According to Theorem 3, an example of our CoarsenRank (Equation (12)) instantiated using PL model can be represented as follows:

$$\text{Equation (12)} \propto \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N P_{\theta}^{\tau_N}(\rho_n) = \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N \left[\prod_{i=1}^{k-1} \frac{\theta_{\rho_n^i}}{\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}} \right]^{\tau_N}. \quad (20)$$

where $\tau_N = \frac{1/N}{1/N+1/\alpha}$. k denotes the length of each preference, which could be variant for different preferences.

Remark 6 (Optimization Intractability and The Data Augmentation Method) *The main inferential issue related to Equation (20) concerns the presence of the annoying normalization terms $\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}, \forall n = 1, 2, \dots, N, i = 1, 2, \dots, k - 1$, that do not permit the direct maximization of the posterior. Further, the nonnegative constraint over the model parameters θ rules out the direct applications of gradient-based optimization approaches.*

Following our analysis in Remark 5, we avoid this issue using the data augmentation method. In particular, we introduce an auxiliary Gamma random variable for each normalization term $\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}, \forall n = 1, 2, \dots, N, i = 1, 2, \dots, k - 1$. Then, the resultant joint distribution would no longer suffers from the annoying normalization terms. ■

3.3 Connection Between CoarsenRank and Mallows Model

Permutation-based models are based on the definition of distances $D(\cdot, \cdot)$, which express the distance of a permutation to the ground truth permutation. The most prominent example of these models is the Mallows model (MM) (Mallows, 1957), an exponential model that expresses the probability of a permutation in terms of its distance to a reference permutation. Let ρ be a full/partial ranking list, then MM specifies:

$$P(\rho|r) = \frac{1}{\psi(\alpha)} \exp(-\alpha D(\rho, r)), \quad \text{where } \psi(\alpha) = \sum_{\rho} \exp(-\alpha D(\rho, r)). \quad (21)$$

Here $\alpha \in \mathbb{R}_+$ is a spread parameter and r is the reference permutation, or called the unknown ground truth. $D(\rho, r)$ represents a sample-level distance between ρ and r . Note that r is the mode, and the closer a ranking list ρ is to r , the larger $p(\rho)$ is. The alternative distance measures considered are Kendall tau distance, Spearman's rank distance, etc (Diaconis, 1988).

According to Equation (21), the probability of MM for \mathcal{R}_N could be represented as

$$\begin{aligned} P(\mathcal{R}_N|r) &= \frac{1}{\psi^N(\alpha)} \prod_{n=1}^N \exp(-\alpha D(\rho_n, r)) = \frac{1}{\psi^N(\alpha)} \exp(-\alpha \sum_{n=1}^N D(\rho_n, r)) \\ &\stackrel{\textcircled{1}}{=} \frac{1}{\psi^N(\alpha)} \exp(-\alpha D(\mathcal{R}_N, r)) \stackrel{\textcircled{2}}{\propto} \exp(-\alpha D(\mathcal{R}_N, r)), \end{aligned} \quad (22)$$

where $\mathcal{R}_N = \{\rho_1, \rho_2, \dots, \rho_N\}$ denotes the collected preferences. $\textcircled{1}$ is valid since we define $D(\mathcal{R}_N, r) = \sum_{n=1}^N D(\rho_n, r)$. $\textcircled{2}$ holds due to the omission of the data non-relevant normalization term.

However, permutation-based models are often impractical for the large-scale problem, because: (1) the normalization term $\psi(\alpha)$ usually requires high computational cost due to discrete distance computation; and (2) a maximum likelihood estimation involves an impossible discrete search for ranking over a large volume of items.

Remark 7 (Comparison Between CoarsenRank and MM) *Comparing Equation (22) to our CoarsenRank formulation (Equation (10)), we can find that the difference between CoarsenRank and MM mainly lies in the definition of distance $D(\cdot, \cdot)$. Namely, distribution-level distance, i.e., $\mathcal{D}_{\text{KL}}(\cdot, \cdot)$ is adopted for CoarsenRank, while sample-level distance, e.g., Kendall tau distance, is adopted for MM. The superiority of CoarsenRank over MM lies in three aspects: (1) In terms of inference, an efficient*

inference strategy is discussed in the Remark after each variant of CoarsenRank, respectively; (2) In terms of explanation, our CoarsenRank formulation is derived from the Coarsening mechanism while assigning an exponential prior for the size of the neighborhood; (3) In terms of optimization w.r.t. the hyperparameter α , we avoid parameter turning by adopting data-driving strategy for choosing α , while Mallow's mode does not enjoy this convenience.

4. Efficient Bayesian Inference

Following the discussion in the previous section, we propose two equivalent algorithms to solve the proposed CoarsenRank framework efficiently, namely a closed-form Expectation-Maximization algorithm in Section 4.2 and a Gibbs Sampling algorithm in Section 4.3. Further, we discuss their algorithm complexity in Section 4.4. To avoid hyperparameter tuning, an efficient model selection method based on Algorithm 2 is introduced in Section 4.5 to choose the single hyperparameter in a data-driven manner.

In particular, we focus on Coarsened PL model (Equation (20)) only and refer it as CoarsenRank for two reasons: (1) PL model can be applied to preferences with various length and incorporates BT model as a special case; (2) Thurstone model and BT model are constrained to pairwise preferences, while the rank breaking strategy would lead to computational inefficient.

4.1 Data Augmentation Method for Eliminating The Normalization Terms

First, we reformulate Equation (20) as follows:

$$\begin{aligned} \text{Equation (20)} &= \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N \left[\prod_{i=1}^{k-1} \frac{\theta_{\rho_n^i}}{\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}} \right]^{\tau N} \\ &= \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N \prod_{i=1}^{k-1} \left[\frac{\theta_{\rho_n^i}}{\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}} \right]^{\tau N} = \max_{\theta \in \Theta} \pi(\theta) \prod_{n=1}^N \prod_{i=1}^{k-1} \left(\frac{\theta_{\rho_n^i}}{\eta_n^i} \right)^{\tau N}, \end{aligned} \quad (23)$$

where $\eta_n^i = \theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}$, $\forall n = 1, 2, \dots, N, i = 1, 2, \dots, k - 1$. Due to the presence of the normalization terms η_n^i in Equation (23), the direct maximization of the posterior w.r.t. to the nonnegative parameter θ would encounter significant inefficiency.

For example, if we want to eliminate the term η_n^i appearing in the denominator, we could introduce an auxiliary random variable, whose conditional distribution contains η_n^i in the numerator. Then, the resultant joint distribution would no longer contain this annoying η_n^i in the denominator. In particular, one particular Gamma distribution is a promising candidate that satisfies the above requirements, i.e., $\text{Gam}(\xi|1, q) = qe^{-q\xi}$ could be used to eliminate a q containing in the denominator.

Therefore, we introduce an auxiliary variable ξ_n^i with regarding to each η_n^i , $\forall n = 1, 2, \dots, N$ and $\forall i = 1, 2, \dots, k - 1$. Further, we define the posterior distribution of ξ_n^i as follows,

$$P(\xi_n^i | \rho_n, \theta) = \text{Gam}(\xi_n^i | 1, \eta_n^i) = \eta_n^i e^{-\xi_n^i \eta_n^i}. \quad (24)$$

Then, we can deal with the joint distribution directly, i.e.,

$$P(\rho_n, \Xi | \theta) = P_{\theta}(\rho_n) \prod_{i=1}^{k-1} P(\xi_n^i | \rho_n, \theta) = \prod_{i=1}^{k-1} \left(\frac{\theta_{\rho_n^i}}{\eta_n^i} \cdot \eta_n^i e^{-\xi_n^i \eta_n^i} \right) = \prod_{i=1}^{k-1} \left(\theta_{\rho_n^i} e^{-\xi_n^i \eta_n^i} \right), \quad (25)$$

Algorithm 1 Closed Form EM for Coarsened Rank Aggregation (CoarsenRank)

- 1: **Input:** the collection of preferences \mathcal{R}_N , the number of iteration T .
 - 2: **Initialization:** hyperparameters $\{a_m, b_m\}_{m=1}^M$ for θ .
 - 3: **for** $t = 1, 2, \dots, T$ **do**
 - 4: *E-step:* calculate the posterior expectation of auxiliary variable ξ_n^i according to Equation (27).
 - 5: *M-step:* update θ_m according to Equation (29) $\forall m = 1, 2, \dots, M$.
 - 6: **end for**
 - 7: **Rank:** the item score θ to derive the final ranking.
 - 8: **Output:** the final ranking.
-

where $\Xi = \{\xi_n^i\}_{i=1}^{k-1}$ denotes the introduced auxiliary variables.

Further, we adopt a Gamma prior to instantiate the prior distribution $\pi(\theta)$, which naturally satisfies the nonnegative constraint of θ , i.e., $\theta \sim \text{Gam}(\theta|a, b) = \prod_{m=1}^M \text{Gam}(\theta_m|a_m, b_m)$. Therefore, the full likelihood of our CoarsenRank model (Equation (23)) can be formulated as follows,

$$\begin{aligned}
 P(\mathcal{R}_N, \Xi, \theta, \epsilon|\{a_m, b_m\}_{m=1}^M, \alpha) &= \pi(\theta) \prod_{n=1}^N \left(P(\rho_n, \Xi|\theta) \right)^{\tau_N} \\
 &= \prod_{m=1}^M \text{Gam}(\theta_m|a_m, b_m) \prod_{n=1}^N \prod_{i=1}^{k-1} \left(\theta_{\rho_n^i} e^{-\xi_n^i \eta_n^i} \right)^{\tau_N},
 \end{aligned} \tag{26}$$

where $\tau_N = \frac{1/N}{1/N+1/\alpha}$. $\mathcal{R}_N = \{\rho_1, \rho_2, \dots, \rho_n\}$ denotes the observed preferences. $\epsilon \sim \text{Exp}(\alpha)$ is the discrepancy between the collected preferences \mathcal{R}_N and its idealized counterpart \mathfrak{R}_N , measured in relative entropy. (a_m, b_m) is initialized to $(1, 2)$, $\forall m = 1, 2, \dots, M$. We fixed $\{a_m, b_m\}_{m=1}^M$ in this paper to eliminate their coupling effects with other factors in CoarsenRank (Equation (26)).

4.2 EM Algorithm with Closed-formed Updating Rules

Concerning the presence of the introduced auxiliary variables Ξ , we resort to the Expectation-Maximization (EM) framework, which is a silver bullet to compute the maximum-likelihood solution or maximum a posterior estimation in the presence of latent variables.

Expectation step (E-step) In the expectation step, we calculate the expectation of each auxiliary variable ξ_n^i w.r.t. its posterior distribution $P(\xi_n^i|\rho_n, \theta)$:

$$\mathbb{E}_{P(\xi_n^i|\rho_n, \theta)}[\xi_n^i] = \frac{1}{\theta_{\rho_n^i} + \theta_{\rho_n^{i+1}} + \dots + \theta_{\rho_n^k}} = \frac{1}{\eta_n^i}, \tag{27}$$

where $n = 1, 2, \dots, N$ and $i = 1, 2, \dots, k-1$. Then, the expectation of the complete-data log-likelihood function w.r.t. the posterior of the introduced auxiliary variables Ξ can be represented as follows:

$$\begin{aligned}
 &\mathbb{E}_{P(\Xi|\mathcal{R}_N, \theta)} \left[\log P(\mathcal{R}_N, \Xi, \theta, \epsilon|\{a_m, b_m\}_{m=1}^M, \alpha) \right] \\
 &= \sum_{m=1}^M \left[(a_m - 1) \log \theta_m - b_m \theta_m \right] + \sum_{n=1}^N \sum_{i=1}^{k-1} \left[\tau_N \log \theta_{\rho_n^i} - \tau_N \mathbb{E}[\xi_n^i] \eta_n^i \right] + \text{constant},
 \end{aligned} \tag{28}$$

where $P(\Xi|\mathcal{R}_N, \theta) = \prod_{n=1}^N \prod_{i=1}^{k-1} P(\xi_n^i|\rho_n, \theta)$ following Equation (27).

Algorithm 2 Gibbs Sampling for Coarsened Rank Aggregation (CoarsenRank)

- 1: **Input:** the collection of preferences \mathcal{R}_N , the number of iteration T .
 - 2: **Initialization:** hyperparameters $\{a_m, b_m\}_{m=1}^M$ for θ , burn-in iterations $\Delta = 100$.
 - 3: **for** $t = 1, 2, \dots, T + \Delta$ **do**
 - 4: sample ξ_n^i from $P(\xi_n^i | \rho_n, \theta^{t-1})$ according to Equation (30), $\forall n = 1, 2, \dots, N, \forall i = 1, 2, \dots, k$.
 - 5: sample θ_m^t from $P(\theta_m | \mathcal{R}_N, \Xi)$ according to Equation (31), $\forall m = 1, 2, \dots, M$.
 - 6: **end for**
 - 7: **Average:** the resulting $\theta^{\Delta+1:\Delta+T}$ to $\bar{\theta}$.
 - 8: **Rank:** the item score $\bar{\theta}$ to derive the final ranking.
 - 9: **Output:** the final ranking.
-

Maximization step (M-step) In the maximization step, we maximize the objective function Equation (28) by setting its gradient w.r.t. θ_m to zero and obtain the following estimates for $\theta_m \forall m = 1, 2, \dots, M$:

$$\theta_m = \frac{\tau_N \sum_{n=1}^N \sum_{i=1}^{k-1} (\varphi_{n,i}^m) + a_m - 1}{\tau_N \sum_{n=1}^N \sum_{i=1}^{k-1} (\psi_{n,i}^m \cdot \mathbb{E}[\xi_n^i]) + b_m}, \quad (29)$$

where $\varphi_{n,i}^m = \begin{cases} 1 & \rho_n^i = m \\ 0 & \text{otherwise} \end{cases}$ and $\psi_{n,i}^m = \begin{cases} 1 & m \in \{\rho_n^i, \dots, \rho_n^k\} \\ 0 & \text{otherwise} \end{cases}$.

Overall, the EM algorithm for Coarsened rank aggregation (CoarsenRank) is summarized in Algorithm 1.

4.3 Gibbs Sampling for Coarsened Rank Aggregation

In our CoarsenRank (Equation (26)), there are two types of latent variables, i.e., Ξ and θ . According to our definition, the posterior distribution of ξ_n^i can be represented as

$$P(\xi_n^i | \rho_n, \theta) = \text{Gam}(\xi_n^i | 1, \eta_n^i) = \eta_n^i e^{-\xi_n^i \eta_n^i}, \quad (30)$$

where $n = 1, 2, \dots, N$ and $i = 1, 2, \dots, k$. Similarly, the full conditional distributions of $\theta_m \forall m = 1, 2, \dots, M$ are still members of the Gamma family. According to Equation (29), the posterior distribution $P(\theta_m | \mathcal{R}_N, \Xi)$ can be represented as

$$P(\theta_m | \mathcal{R}_N, \Xi) = \text{Gam} \left(\theta_m \mid \tau_N \sum_{n=1}^N \sum_{i=1}^{k-1} (\varphi_{n,i}^m) + a_m - 1, \tau_N \sum_{n=1}^N \sum_{i=1}^{k-1} (\psi_{n,i}^m \cdot \mathbb{E}[\xi_n^i]) + b_m \right), \quad (31)$$

where $\varphi_{n,i}^m = \begin{cases} 1 & \rho_n^i = m \\ 0 & \text{otherwise} \end{cases}$ and $\psi_{n,i}^m = \begin{cases} 1 & m \in \{\rho_n^i, \dots, \rho_n^k\} \\ 0 & \text{otherwise} \end{cases}$.

Therefore, the Gibbs sampling procedure for Coarsened rank aggregation (CoarsenRank) can be summarized in Algorithm 2.

4.4 Time and Space Complexity Analysis

We analyze the algorithm complexity for our CoarsenRank as well as vanilla RA methods to further demonstrate the superiority of our CoarsenRank. Note that CoarsenBT and CoarsenPL denote the

Complexity		Optimization	Time	Space
Vanilla	TH	Gradient descent	$O(NMT_1k^2)$	$O(NMk^2)$
	BT	Expectation maximization	$O(NMT_2k^2)$	$O(NMk^2)$
	PL-EM	Expectation maximization	$O(NMT_2k^2)$	$O(NMk)$
	PL	Gradient descent	$O(NMT_1k^2)$	$O(NMk)$
Robust	CrowdBT	Bayesian moment matching	$O(NT_1k^2)$	$O(\max(M, G))$
	PeerGrader	Gradient descent	$O(NMT_1k^2)$	$O(NMk)$
Coarsen	CoarsenTH	Gradient descent	$O(NMT_1k^2)$	$O(NMk^2)$
	CoarsenBT	Expectation maximization	$O(NMT_2k^2)$	$O(NMk^2)$
	CoarsenPL	Expectation maximization	$O(NMT_2k^2)$	$O(NMk)$
	CoarsenPL(GS)	Gibbs sampling	$O(NM(T_3 + \Delta)k^2)$	$O(NMk + MT_3)$

Table 2: Complexity analysis of CoarsenRank

optimization strategy in Algorithm 1, while CoarsenPL(GS) denotes the optimization strategy in Algorithm 2.

Let N , M , T_1 , T_2 , and T_3 denote the number of preferences, the number of items, the number of gradient descent steps, the number of EM iterations, and the number of Gibbs sampling iterations, respectively. For ease of analysis, we assume the length of all preferences equals k . The time and space complexities of CoarsenRank are listed in Table 2.

Let’s give more discussions about Algorithm 1 and Algorithm 2 in terms of time complexity:

1. Our CoarsenRank consists of T iterations between Equation (27) and Equation (29). The main computational routine is the denominator $\sum_{n=1}^N \sum_{i=1}^{k-1} (\psi_{n,i}^m \cdot \mathbb{E}[\xi_n^i])$ for $m = 1, 2, \dots, M$, whose computational cost is $O(NMk^2)$ per iteration. The same computations also apply to PL-EM since the only difference between PL-EM and CoarsenPL is the value of the scalar τ_N , which refers to PL-EM if $\tau_N = 1$, and CoarsenPL if $0 < \tau_N < 1$.
2. Pairwise RA methods, e.g., CoarsenTH and CoarsenBT, need to split each k -ary preference into $\frac{k(k-1)}{2}$ pairwise preferences, thus increasing the number of preferences from N to $\frac{Nk(k-1)}{2}$, but shorten the computation cost per preference from $O(\frac{k(k-1)}{2})$ to $O(1)$. The analysis also apply to TH and BT since the only difference between CoarsenRank (CoarsenTH and CoarsenBT) and vanilla RA (TH and BT) is the value of the scalar τ_N .
3. As for each iteration, gradient descent-based optimization (TH, CoarsenTH and PL) has the same computational cost, i.e., $O(NMk^2)$ as that of EM-based ones (BT, CoarsenBT, PL-EM and CoarsenPL), since it needs to calculate the gradient of M items over $\frac{Nk(k-1)}{2}$ pairwise preferences for BT-based models or N k -ary preferences for PL-based models.
4. CrowdBT adopts Bayesian moment matching for online updates. Although with complex updating rules, it only updates the emerging item in the current preference without traversing all items. Namely, its computation cost is not relevant to M .

5. PeerGrader introduces an extra variable to model the grader reliability. Since it adopts gradient descent for optimization, PeerGrader has the same time complexity as PL. However, an iterative alternating-minimization is required for optimizing the item score and the grader reliability, which incurs more computations than vanilla PL.
6. The analysis for CoarsenPL in Algorithm 1 also applies to CoarsenPL(GS) of Algorithm 2 by comparing Eq.31 and Eq.29, which, however, requires more iterations (i.e., $T_3 \gg T_2$). Furthermore, CoarsenPL(GS) also includes an extra computation cost $O(NM\Delta k^2)$ for the burn-in process.

In terms of space complexity:

1. The primary storage elements are the index $\varphi_{n,i}^m$ and $\psi_{n,i}^m$, where $n = 1, 2, \dots, N, i = 1, 2, \dots, k$ and $m = 1, 2, \dots, M$. Therefore, the overall space complexity for PL-based methods, e.g., PL, PL-EM and CoarsenPL, is $O(NMk)$.
2. Apart from the same storage as other PL-based methods, PeerGarder needs to store the reliability (a scalar) for each grader as well, which however can be omitted compared to the storage of indices. Therefore, the space complexity for PeerGarder is also $O(NMk)$.
3. Pairwise RA methods, e.g., TH, CoarsenTH and BT, need to split each k -ary preference into $\frac{k(k-1)}{2}$ pairwise preferences, thus increasing the number of preferences from N to $\frac{Nk(k-1)}{2}$, so the overall space complexity is $O(NMk^2)$.
4. Since CrowdBT adopts an online updating paradigm, it only needs to store a two-dimensional hyperparameter for each score, a two-dimensional hyperparameter for each grader. Let G denote the number of graders, the overall space complexity for CrowdBT is $O(\max(M, G))$.
5. Similarly, apart from the storage of the two index matrices φ and ψ , CoarsenPL(GS) requires extra storage about MT_3 for storing T_3 intermediate variables ($\theta \times T_3$) so as to calculate the final item score.

Note that each item only appears in a small number of preferences, which means there is no need to store all N preferences for each item, especially when the length of preference (k) is small. Therefore, we can choose a more efficient strategy for storing the preferences, namely, regarding each item, we only store the relevant preferences that contain the target item. Assume the maximum number of relevant preferences is $\tilde{N} \ll N$, which can reduce both the time complexity and space complexity of CoarsenPL from $O(NMT_2k^2)$ and $O(NMk)$ to $O(\tilde{N}MT_2k^2)$ and $O(\tilde{N}Mk)$, respectively. The same is true when applying to other RA methods. We adopt the proposed strategy for all RA methods in the experiment.

4.5 A Data-driven Strategy for Choosing α

Regarding the hyperparameter optimization, we focus on exploring the effects of the hyperparameter α in this paper. Other hyperparameters were not further explored in our experiment since these are not our focus in this paper.

Regarding the hyperparameter α , we have no prior basis for choosing parameter α in Equation (10). Therefore, the following diagnostic curve can help to make a data-driven choice. Let $f(\alpha)$ be a measure of fit to the data and $g(\alpha)$ be a measure of model complexity. Following (Spiegelhalter

Baselines		Distance measure	Perturbation assumption		Ranking model assumption		
		distribution or sample level	perturbation $P(\rho_n \mathcal{Q}_n)$	neighborhood size ϵ	model P_θ	prior $\pi(\theta)$	posterior $P(\xi_n^i \rho_n, \theta)$
Vanilla	TH	distribution	—	—	TH	Gaussian	Gaussian
	BT	distribution	—	—	BT	Gaussian/Gamma	same as prior
	PL	distribution	—	—	PL	Gaussian/Gamma	same as prior
Robust	MM	sample	fractional likelihood	—	MM	—	—
	CrowdBT	distribution	Dawid-Skene model	—	BT	Gaussian	Gaussian
	ROPAL	distribution	Dawid-Skene model	—	PL	Gaussian	Gaussian
	PeerGrader	distribution	fractional likelihood	—	BT/PL/TH	Gaussian	—
Coarsen	CoarsenTH	distribution	—	exponential prior	TH	Gaussian	—
	CoarsenBT	distribution	—	exponential prior	BT	Gamma	Gamma
	CoarsenPL	distribution	—	exponential prior	PL	Gamma	Gamma

Table 3: Comparison between various ranking models in terms of their distance measure, perturbation assumption and model assumption. “—” denotes the assumption does not apply to a particular ranking model.

et al., 1998), we use the posterior expected log-likelihood for $f(\alpha)$, and the difference between the log-likelihood evaluated at the posterior mean of the parameters and the posterior expected log-likelihood for $g(\alpha)$. Specifically, we define

$$f(\alpha) = \mathbb{E}_{q_\alpha(\theta|\mathcal{R}_N)} [\log P_\theta(\mathcal{R}_N)] \quad \text{and} \quad g(\alpha) = \log P_{\mathbb{E}(\theta)}(\mathcal{R}_N) - f(\alpha), \quad (32)$$

where $q_\alpha(\theta|\mathcal{R}_N)$ is an approximate posterior distribution for θ , i.e., Equation (31), and $\mathbb{E}(\theta) = \mathbb{E}_{q_\alpha(\theta|\mathcal{R}_N)}[\theta]$ is the posterior expectation of θ .

Therefore, the adopted Deviance Information Criterion (DIC) (Spiegelhalter et al., 1998) can be represented as

$$\text{DIC} = g(\alpha) - f(\alpha). \quad (33)$$

As α ranges from 0 to $+\infty$, DIC traces out a curve in \mathbb{R} , and the technique is to choose α with the lowest DIC or where DIC levels off.

5. Perturbation Assumption and Rank Model Assumption

In this section, we summarize the differences between our CoarsenRank and related RA models in terms of distance measure, perturbation assumption and ranking model assumption. A detailed comparison is listed in Table 3.

Specifically, we classify the rank aggregation methods into three categories: Vanilla RA, Robust RA, and Coarsen RA. Vanilla RA, such as Thurstone (TH) model (Thurstone, 1927b) Bradley-Terry (BT) model (Bradley and Terry, 1952) Plackett-Luce (PL) model (Plackett, 1975; Luce, 1959), assumes all ranking lists are generated from the same distribution, under some parameterized ranking model, e.g., TH, BT, and PL. Further, a prior is usually introduced for the parameter to avoid overfitting.

The formulation of MM is similar to the fractional likelihood (Bhattacharya et al., 2019) where the tunable spread parameter in MM has the same function as the exponential variable in a fractional likelihood. This is why we classify MM as a robust RA. See section 3.3 for more detailed discussion.

Robust RA extends vanilla RA by considering extra noisy perturbations incurred during the data collection. Representative robust RAs are CrowdBT (Chen et al., 2013) and ROPAL (Han et al., 2018), which recovers the idealized ranking lists by some predefined perturbation mechanisms, following the Dawid-Skene model. Meanwhile, PeerGrader (Raman and Joachims, 2014) also enhances the robustness of vanilla RA following the principle of fractional likelihood. In particular, PeerGrader introduces an extra parameter for each user, which would decrease the effects of noisy preferences from an unreliable user. Above Robust RAs aim at capturing a perturbation pattern w.r.t. to each user, which requires sufficient available preferences from each user. This constraint is too strong while only one preference from each user is usually available in real application.

Our CoarsenRank is motivated by the Coarsening mechanism, where we perform regular RA over the neighborhood of original ranking data. We only introduce one extra parameter, representing the size of the neighborhood. A simple formulation of CoarsenRank is further derived if we adopt relative entropy as the distance measure and assign an exponential prior for the unknown neighborhood size. Three variants of CoarsenRank are introduced by instantiating with the vanilla TH, BT, and PL model, and an efficient EM algorithm could be derived under the same assumption as the vanilla RAs.

6. Experimental Evaluation

In this section, we verify the efficacy of the proposed CoarsenRank algorithm on noisy rank aggregation with the state-of-the-art approaches. The results are carried on four real-world noisy ranking data sets.

6.1 Experimental Setting

Performance metric: Regarding the performance metric, we consider the Kendall tau similarity (Kendall, 1938), which is one of the most common measures of similarity between rankings, namely

$$\tau(\sigma_1, \sigma_2) = \frac{1}{M} \sum_{i < j} (\mathbb{1}[(\sigma_1^i > \sigma_1^j) \wedge (\sigma_2^i > \sigma_2^j)] + \mathbb{1}[(\sigma_1^i < \sigma_1^j) \wedge (\sigma_2^i < \sigma_2^j)]). \quad (34)$$

$\tau(\sigma_1, \sigma_2)$ counts the pairwise agreements between items from two rankings σ_1 and σ_2 . $\bar{M} = \frac{1}{2}M(M - 1)$ denotes total number of pairs. τ ranges from 0 (worst) to 1 (best).

Baselines: As for vanilla baselines, we first consider the vanilla Plackett-Luce model (Plackett, 1975; Luce, 1959). For the sake of fair comparison, we optimize with two optimization approaches, i.e., gradient descent (PL) (Boyd and Vandenberghe, 2004) and EM using data augmentation (PL-EM) (Caron and Doucet, 2012).

In terms of robust RA, we compare the results with PeerGrader (Raman and Joachims, 2014), which is a variation of the Plackett-Luce model for partial preferences while incorporating the user reliability estimation module. We also compare with the popular noisy ranking model CrowdBT (Chen et al., 2013). Since CrowdBT was originally designed for pairwise preferences, we generalize CrowdBT to partial preferences following rank-breaking (Weng and Lin, 2011). Namely, we first break each partial preference into a set of pairwise comparisons and then apply CrowdBT to each pairwise comparison independently. Note that we do not consider the robust RA proposed by Han et al. (2018), since it is tailor-design for the crowdsourcing setting. In particular, their method requires multiple preferences from each user for initializing the parameters, while it is not true for a more general setting considered in this work where only one preference from each user is usually available.

In terms of our CoarsenRank, we consider all its three instantiations: CoarsenTH (Equation (16)), CoarsenBT (Equation (18)) and CoarsenPL (Equation (20)). Similar to CrowdBT, we adopt the rank-breaking to generalize the CoarsenTH and CoarsenBT to partial preferences.

Calibration for real application In real applications, the number of items involved in partial comparisons usually varies significantly. Some items (e.g., user-friendly cameras, or popular sushi) may appear frequently in the ranking list due to their popularity, while other items (e.g., professional cameras, or sushi with special taste) appear only in few preferences due to their professionalism. In such cases, the final ranking will not be unique or even not converge. To ensure a unique solution and to avoid overfitting, regularization may be used. Specifically, we perform normalization over θ , i.e., $\theta = C \cdot \theta / \sum_m \theta_m$. We fixed $C = M/2$ in our experiment for simplicity. The calibration method is applied to all EM-based approaches, i.e., CoarsenBT, CoarsenPL, and PL-EM. In terms of other baselines, e.g., PL, CrowdBT and CoarsenTH, their formulation is a little different and the nonnegative constraint is no longer required. Therefore, the calibration method cannot be applied. Following CrowdBT, we use virtual node regularization (Chen et al., 2013). Specifically, it augments the original data set \mathcal{R}_N with \mathcal{R}_o , which consists of the pairwise comparisons between all items and a virtual item θ_0 , namely $\mathcal{R}_o = \{\theta_m > \theta_0, \theta_m < \theta_0, m = 1, 2, \dots, M\}$.

6.2 Detailed Descriptions of Data sets

We conducted our experiment on four real-world data sets introduced in previous research. The detailed descriptions of the data sets are introduced in the following.

The *Readlevel* data set (Chen et al., 2013) contains English text excerpts whose reading difficulty level is annotated by workers from a crowdsourcing platform. This data set consists of 490 excerpts from 624 workers, resulting in a total of 12,728 pairwise comparisons. A total order for 490 excerpts provided by the domain expert is regarded as ground truth. The number of annotations varies significantly for different workers, ranging from 1 to 42.

The *SUSHI* data set is introduced in (Kamishima, 2003), which consists of partial preferences over 100 types of sushi from $n = 5,000$ customers. Following (Khetan and Oh, 2016), we generated the total order as ground truth using the vanilla PL over the entire 5,000 preferences. To create training data, we randomly replaced 2,000 preferences in the original *SUSHI* data set with another 2,000 random generated preferences. The number of preference provided by each customer is fixed to one.

The *BabyFace* data set (Han et al., 2018) consists of the evaluations of workers from Amazon Mechanical Turk on images of children’s facial microexpressions from happy to angry, which yields a collection of 3,074 trinary preferences from 41 workers. A total order over all microexpressions is provided as ground truth by the agreement of most workers after the experiment. Each worker provides at least 60 annotations.

The *PeerGrading* data set (Sajjadi et al., 2016) consists of assessments, i.e., Self grading and Peer grading, from 219 students over 79 group submissions. We then created the ordinal gradings by merging the Self grading and Peer grading regarding the same assignment provided by each student, which results in a total of 3,619 preferences with each containing 2 or 3 items. Further, the TA gradings (following a linear order) provided by six teaching assistants over all submissions are considered as ground truth. The number of annotations from different students ranges from 2 to 26.

The statistics of four data sets are summarized in Table 4.

Data set	#items (M)	#users	#preferences (N)	length of preferences (k)	#annotations per user
Readlevel	490	624	12,728	2	1 – 42
SUSHI	100	5,000	5,000	10	1
BabyFace	18	41	3,074	3	≥ 60
PeerGrading	219	79	3,619	2 or 3	2 – 26

Table 4: The statistics of four real ranking data sets

6.3 Exploring The Efficacy of The Calibration Method

In section 6.1, we introduce a calibration method to the vanilla EM algorithm to deal with data sparsity. We claim that the calibration method is necessary when the collected preferences do not evenly cover the items. In particular, the calibration method serves as regularization, which is helpful to avoid overfitting and leads to a unique solution.

To verify the efficacy of the calibration method, we conducted rank aggregation using CoarsenPL and PL-EM on the *Readlevel* data set. Then we collected the Kendall tau similarity of these two methods in Figure 3, under the case of with or without calibration method, respectively.

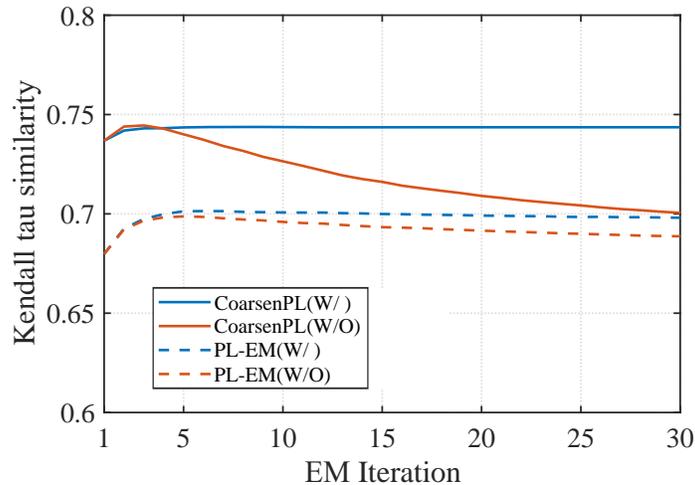


Figure 3: The performance comparison of CoarsenRank and PL-EM algorithm under the case of with or without calibration method. “W/” denotes “with” while “W/O” denotes “without”.

Figure 3 shows that: (1) without the calibration method, CoarsenPL and PL-EM are all prone to overfitting. In particular, the Kendall tau similarity of CoarsenPL and PL-EM can reach their optima at the first few iterations but start to decrease in the later iterations. Since the Kendall tau similarity is different from our objective, this phenomenon is considered as a sign of overfitting. (2) With the help of the calibration method, the overfitting problem is avoided. The Kendall tau similarity of two methods remains stable after reaching their optima, respectively. (3) With or without the calibration, the optimum Kendall tau similarities of CoarsenPL are very close. The same is true for PL-EM. It implies that the calibration method does not change the result of rank aggregation methods but just avoids overfitting.

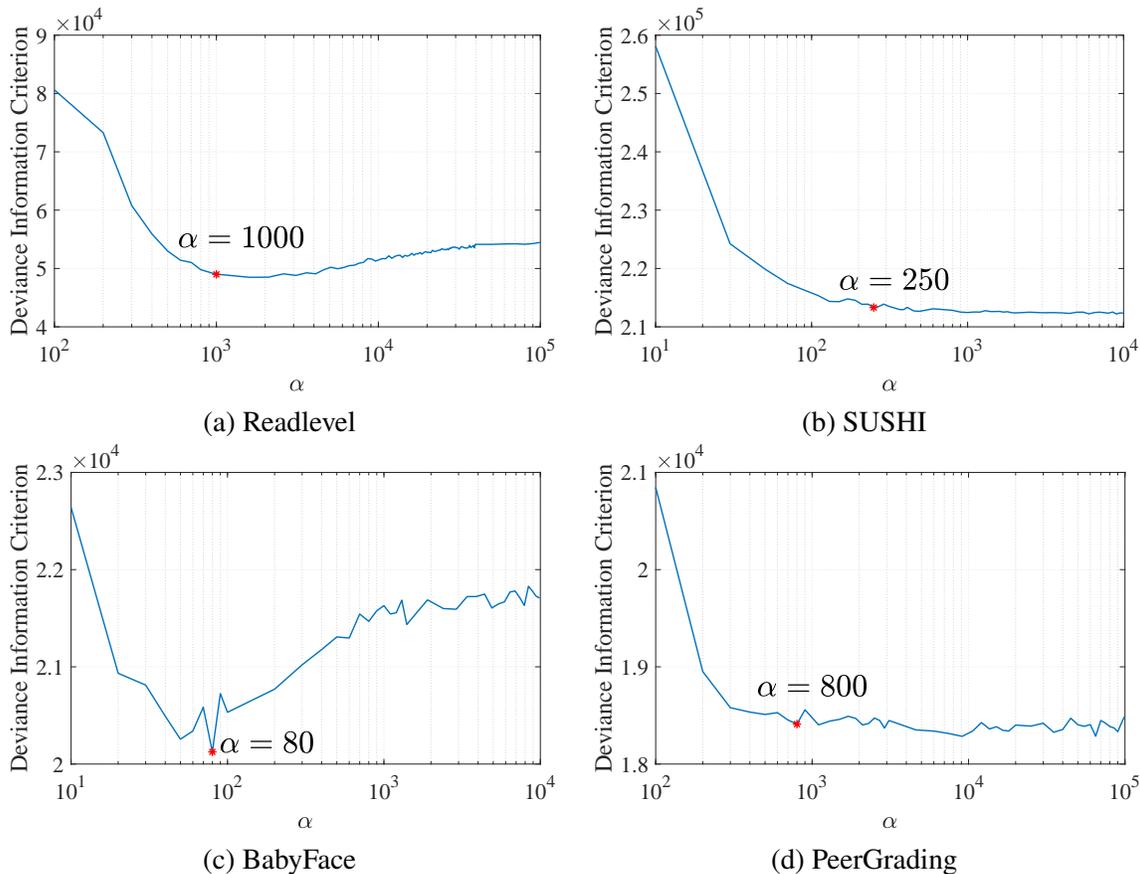


Figure 4: (a)-(d) The diagnostic plot of DIC VS. α on four data sets, respectively. The α used in our experiment are marked as “*” in each figure.

6.4 Deviance Information Criterion (DIC) for Choosing The Hyperparameter α

Following Section 4.5, we adopted the DIC to choose the hyperparameter α for different data sets. Since it is intractable to analytically calculate the posterior expectation in DIC, we implemented a Gibbs Sampling procedure in Algorithm 2. Then, we collected the samplings from $P(\theta_m | \mathcal{R}_N, \Xi)$ (Equation (31)) and calculated the Monte Carlo estimation of DIC (Equation (33)) for different α . The number of samplings is set to 50 in our experiment. The diagnostic curves of α on four data sets are plotted in Figure 4, respectively.

The results show that the DIC decreases dramatically at first when α is small, then the curve reaches a cusp and levels off, with more modest increases/decreases when α becomes larger. α is chosen at the point with the lowest DIC or where DIC levels off in our experiment, marked as “*” in each figure.

6.5 The Accuracy of CoarsenRank in Four Real Applications

We conducted the experiments on four real data sets and collected the empirical results of various rank aggregation methods in Table 5. For better comparison, we further showed the performance improvement of all methods over PL-EM on four data sets in Figure 5.

Baselines		Readlevel	SUSHI	BabyFace	PeerGrading	Ave. Rank (\downarrow)
Vanilla	PL	0.6853	0.8554	0.8824	0.8023	6.75
	PL-EM	0.6879	0.8578	0.8824	0.8014	6.25
Robust	CrowdBT	0.6944	0.8765	0.9085	0.8060	3
	PeerGrader	0.6965	0.8588	0.9150	0.8040	3.5
Coarsen	CoarsenTH	0.6897	0.8731	0.9085	0.8095	3.5
	CoarsenBT	0.7436	0.8740	0.8889	0.8117	2.75
	CoarsenPL	0.7436	0.8970	0.9020	0.8130	1.75

Table 5: Kendall tau similarity (Eq.(34)) of various rank aggregation methods on four real data sets. Best results are marked in bold.

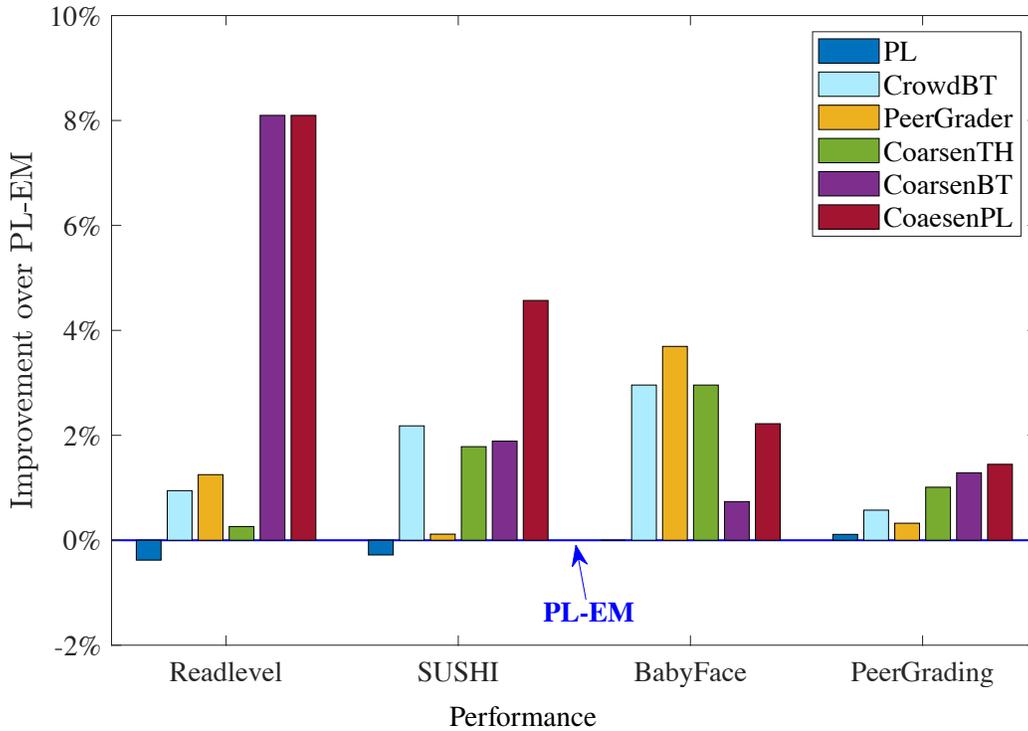


Figure 5: Performance improvement of various methods over PL-EM on four data sets in terms of Kendall tau similarity (Eq.(34)), following $\frac{\tau_s - \tau_0}{\tau_0}$. τ_0 is the Kendall tau similarity of PL-EM.

6.5.1 COMPARISON BETWEEN THREE COARSEN RANK VARIANTS AND OTHER BASELINES

From Table 5 and Figure 5, it can be observed that: (1) Three CoarsenRank variants achieve consistent improvement over other rank aggregation baselines. It demonstrates the great potential of CoarsenRank in real applications, where model misspecification widely exists. (2) The accuracy of PL is comparable with PL-EM on all data sets, which rules out the possibility that the EM algorithm would lead to performance improvement. (3) CrowdBT and PeerGrader get superior performance on *BabyFace* because of sufficient annotations (over 60) from each user and the trinary preferences setting in *BabyFace*. (4) The improvement of CrowdBT and PeerGrader vary significantly on different data sets. The reason is that their pre-assumed perturbation patterns may not be consistent with noise agnostic perturbations in different data sets. (5) Marginal improvement is achieved by CrowdBT and PeerGrader on *Readlevel*, *SUSHI* and *PeerGrading* where each user provides almost one preference. Points (4) & (5) are model misspecification cases which our CoarsenRank is intended to address.

6.5.2 COMPARISON AMONG THREE COARSEN RANK VARIANTS

From the fourth column of Table 5, we can find that: (1) CoarsenPL achieves the highest Kendall tau similarity on most data sets, compared to other CoarsenRank variants. (2) CoarsenBT gets the inferior performance to CoarsenPL. It is because CoarsenBT are tailored designed for pairwise preferences, and it needs to break each partial preference into independent pairwise preferences before aggregation, which has been recently shown to introduce inconsistency (Khetan and Oh, 2016). (3) CoarsenTH achieves the lowest Kendall tau similarity among all CoarsenRank variants. Apart from the above-mentioned inconsistency issue, CoarsenTH may suffer from accuracy reduction due to the introduced approximation or the lack of closed-form updating solutions.

Ratio		1%	2%	4%	8%	16%	32%	64%	100%
Average annotations		1.19	3	6	11.95	23.43	46.76	93.67	146.38
Vanilla	PL	0.7778	0.7582	0.8497	0.8693	0.8824	0.8824	0.8824	0.8824
	PL-EM	0.6928	0.6863	0.8235	0.8497	0.8497	0.8562	0.8824	0.8824
Robust	CrowdBT	0.7843	0.7908	0.8627	0.8824	0.8824	0.9085	0.9020	0.9085
	PeerGrader	0.8497	0.8562	0.8562	0.8824	0.9085	0.9085	0.9150	0.9150
Coarsen	CoarsenTH	0.8824	0.8824	0.8824	0.8824	0.8824	0.8824	0.8889	0.9085
	CoarsenBT	0.8824	0.8824	0.8824	0.8824	0.8824	0.8824	0.8889	0.8889
	CoarsenPL	0.8301	0.8431	0.8627	0.8693	0.8758	0.9085	0.9085	0.9020

Table 6: Ablation study on the number of annotations per user using the *BabyFace* data set by gradually reducing the number of annotations per user. The Kendall tau similarity is calculated for various rank aggregation methods. Best results are marked in bold.

6.5.3 ABLATION STUDY ON THE NUMBER OF ANNOTATIONS PER USER

To justify our claim, we conducted an ablation study on the *BabyFace* data set by gradually reducing the number of annotations per user and collected the results in the following.

Table 6 shows that: (1) the performance of three CoarsenRank variants is less sensitive to the number of annotations compared to other vanilla and robust rank aggregation methods; (2) when the number of annotation per user is insufficient (i.e., ≤ 12) to learn the ability of annotator, three CoarsenRank variants can achieve the best aggregation performance; (3) the performance of CoarsenPL is significantly better than its vanilla version, i.e., PL-EM, which demonstrates the superiority of CoarsenRank for noisy rank aggregation.

6.6 Comparisons of All Methods w.r.t. The Learning Time

We independently ran each baseline 50 times and reported the learning time (s) in Table 7 and Figure 6. The learning time is represented by the mean with the standard deviation. For the sake of fair comparison, the inner iteration is fixed to 15 for all methods. Empirical analyses were performed on an Intel i5 processor(2.30 GHz) and 8 GB random-access memory (RAM).

Table 7 and Figure 6 shows that: (1) Three CoarsenRank variants achieve much smaller learning time compared to other robust ranking aggregation baselines. It shows that our CoarsenRank is promising for deploying in a large-scale environment, where reliability and efficiency are both required. (2) The learning time of CoarsenPL and PL-EM are comparable because of the only difference between CoarsenPL and PL-EM lying at the choosing of parameter τ (See Equation (29)). (3) PeerGrader suffers from significant inefficiencies since it needs to optimize parameters alternatively. (4) CrowdBT replaces the inefficient alternative optimization with the online Bayesian moment matching and achieves smaller learning time compared to PeerGrader. However, it is still time-consuming on *SUSHI* data set because of the inefficient rank-break method for long preferences.

6.6.1 ABLATION STUDY ON THE SIZE OF SAMPLES

To demonstrate the time complexity analysis in Section 4.4, we took the *Readlevel* data set as an example for the experiment since it has the largest number of samples (13k). We first magnified the

Baselines		Readlevel	SUSHI	BabyFace	PeerGrading	Ave. Rank(\downarrow)
Vanilla	PL	0.793 \pm 0.050	10.416 \pm 1.605	1.522 \pm 0.087	1.250 \pm 0.075	4.75
	PL-EM	1.352 \pm 0.124	2.305 \pm 0.338	0.475 \pm 0.038	0.408 \pm 0.037	2.5
Robust	CrowdBT	3.831 \pm 0.232	15.169 \pm 1.598	1.193 \pm 0.0554	0.819 \pm 0.072	5.5
	PeerGrader	36.580 \pm 0.967	57.289 \pm 1.347	27.217 \pm 0.973	23.966 \pm 0.128	7
Coarsen	CoarsenTH	0.253 \pm 0.014	0.706 \pm 0.029	0.050 \pm 0.011	0.117 \pm 0.018	1
	CoarsenBT	1.355 \pm 0.102	9.155 \pm 0.308	0.748 \pm 0.059	0.488 \pm 0.016	4
	CoarsenPL	1.364 \pm 0.098	2.350 \pm 0.671	0.459 \pm 0.012	0.389 \pm 0.024	3.25

Table 7: The learning time (s) of all baselines and the proposed ones on four data sets, respectively. Best results are marked in bold.

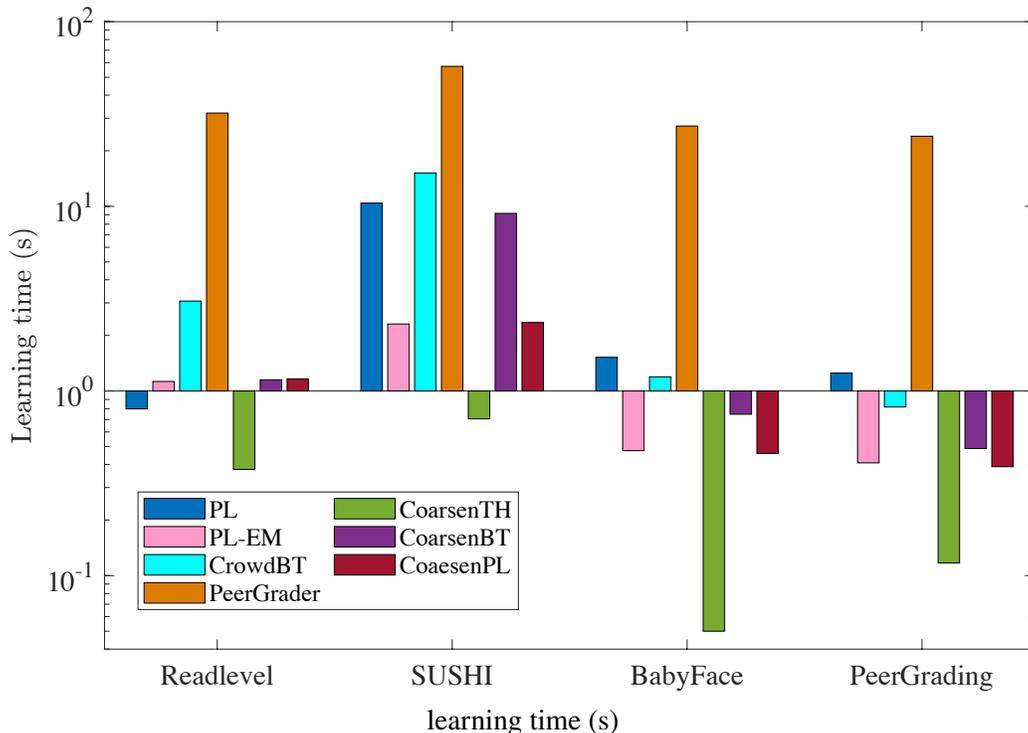


Figure 6: The learning time of all baselines and the proposed ones on four data sets, respectively.

Magnification ratio		1	10	50	100
Vanilla	PL	0.793 ± 0.050	7.445±0.113	37.445± 0.298	72.915±0.960
	PL-EM	1.352 ± 0.124	12.909 ± 0.286	64.081 ± 0.690	124.838 ± 2.222
Robust	CrowdBT	3.831 ± 0.232	38.703 ± 0.573	193.832 ± 1.105	378.552 ± 8.037
	PeerGrader	36.580±0.967	520.240±9.908	3557.707±21.783	7113.333±69.421
Coarsen	CoarsenTH	0.253 ± 0.014	2.367±0.028	12.258±0.213	23.650±0.383
	CoarsenBT	1.355 ± 0.102	12.793 ± 0.212	64.111 ± 0.663	124.851 ± 1.887
	CoarsenPL	1.364 ± 0.098	12.858 ± 0.222	64.176 ± 0.691	125.343 ± 2.704

Table 8: Ablation study on the size of the data using the *Readlevel* data set by gradually magnifying the size of data. The learning time (s) is represented by the mean with the standard deviation. Best results are marked in bold.

Readlevel data set by duplicating it with the specified ratios and then collected the learning time (s) of various methods on each magnified data set, respectively. For the sake of fair comparison, the inner iteration is fixed to 15 for all methods. The learning times are summarized in Table 8 and Figure 7,

From Table 8 and Figure 7, we can find that: (1) the learning time of all the methods exhibit obvious linear correlation with the size of the data set. This is consistent with our analysis in Section 4.4 that the learning time is linearly correlated with the number of samples. (2) The

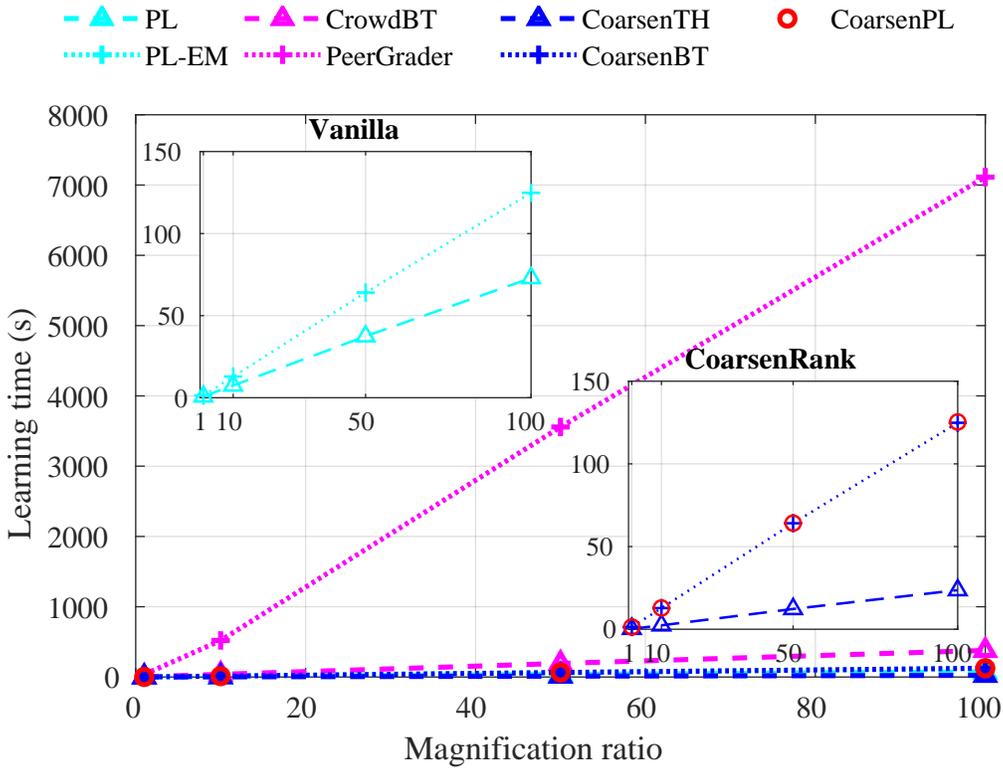


Figure 7: Visualization of the learning time (s) for various rank aggregation methods on the magnified *Readlevel* data set.

superiority of our CoarsenRank variants over PeerGrader is significant. In particular, for a large-scale data set with $100 * 13k$ samples, CoarsenTH can output the result in less than one minute, while PeerGrader needs almost 2 hours to get the result. (3) CoarsenRank (i.e., CoarsenBT and CoarsenPL) brings almost no extra computation cost over its vanilla counterpart (i.e., PL-EM) regardless of the number of samples.

7. Conclusion

Our CoarsenRank performs imprecise inference conditioning on a neighborhood of the ranking data set, which opens a new door to the robust rank aggregation against model misspecification. Particularly, a computationally efficient formula for CoarsenRank is derived, which introduces only one extra hyperparameter to vanilla ranking models. Experiments on four real applications demonstrate imprecise inference on the neighborhood of the preferences, instead of the original data set, can improve the model reliability. It shows that our CoarsenRank has great potential in real applications, e.g., social choice, information retrieval, recommender system, etc, where model misspecification widely exists. We consider the rank aggregation problem where only one ground truth consensus full ranking exists. In terms of other applications for mixture rank aggregation (Zhao et al., 2016), the homogeneity assumption on users is no longer valid and we will extend our distributionally robust CoarsenRank to the scenario of heterogeneity users. Another promising direction for future research is to explore other divergence metrics for other statistical properties of rank aggregation.

Acknowledgments

IWT is supported by ARC under grants DP180100106 and DP200101328. WC is supported by the national NSFC (No.61603338 and No.11426202) and the Zhejiang Provincial NSFC (LY21F030013). GN is supported by JST AIP Acceleration Research Grant Number JPMJCR20U3, Japan. MS is supported by the International Research Center for Neurointelligence (WPI-IRCIN) at The University of Tokyo Institutes for Advanced Study.

Appendix A. Proof for Theorem 1

Theorem 1 Suppose $D(\mathcal{R}_N, \mathfrak{X}_N)$ is an almost surely (a.s.)-consistent estimator⁴ of $D(P_o, P_\theta)$, namely $D(\mathcal{R}_N, \mathfrak{X}_N) \xrightarrow[N \rightarrow +\infty]{\text{a.s.}} D(P_o, P_\theta)$, where $F_N(x|\mathcal{R}_N) \rightarrow P_o$ and $F_N(x|\mathfrak{X}_N) \rightarrow P_\theta$ when $N \rightarrow +\infty$. Assume $\mathbb{P}(D(P_o, P_\theta) = \epsilon) = 0$ and $\mathbb{P}(D(P_o, P_\theta) < \epsilon) > 0$, then we have

$$\mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon) \xrightarrow[N \rightarrow +\infty]{\text{a.s.}} \mathbb{P}(\theta | D(P_o, P_\theta) < \epsilon),$$

for any $\theta \in \Theta$ such that $\int |\theta| \mathbb{P}(d\theta) < \infty$.

Proof: Since $\mathbb{P}(D(P_o, P_\theta) = \epsilon) = 0$, we have⁵ $\mathbb{1}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon) \xrightarrow[N \rightarrow +\infty]{\text{a.s.}} \mathbb{1}(D(P_o, P_\theta) < \epsilon)$. Then we have $|\mathbb{1}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon)| \leq 1$ hold $\forall N > 0$ since $0 \leq \mathbb{1}(\cdot) \leq 1$. According to the dominated convergence theorem (Billingsley, 1999), we have $\mathbb{P}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon) \xrightarrow[N \rightarrow +\infty]{\text{a.s.}} \mathbb{P}(D(P_o, P_\theta) < \epsilon)$.

Similarly, we have $\theta \mathbb{1}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon) \xrightarrow[N \rightarrow +\infty]{\text{a.s.}} \theta \mathbb{1}(D(P_o, P_\theta) < \epsilon)$. Since $0 \leq \mathbb{1}(\cdot) \leq 1$, $|\theta \mathbb{1}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon)| \leq |\theta|$, $\forall N > 0$. Therefore, we have $\mathbb{P}(\theta \mathbb{1}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon)) \xrightarrow[N \rightarrow +\infty]{\text{a.s.}} \mathbb{P}(\theta \mathbb{1}(D(P_o, P_\theta) < \epsilon))$, according to the dominated convergence theorem and $\int |\theta| \mathbb{P}(d\theta) < \infty$. Above all, we have

$$\begin{aligned} \mathbb{P}(\theta | D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon) &= \frac{\mathbb{P}(\theta \mathbb{1}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon))}{\mathbb{P}(D(\mathcal{R}_N, \mathfrak{X}_N) < \epsilon)} \\ &\xrightarrow[N \rightarrow +\infty]{\text{a.s.}} \frac{\mathbb{P}(\theta \mathbb{1}(D(P_o, P_\theta) < \epsilon))}{\mathbb{P}(D(P_o, P_\theta) < \epsilon)} = \mathbb{P}(\theta | D(P_o, P_\theta) < \epsilon). \quad \blacksquare \end{aligned}$$

References

- Soroosh Shafieezadeh Abadeh, Peyman Mohajerin Mohajerin Esfahani, and Daniel Kuhn. Distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, pages 1576–1584, 2015.
- Syed Mumtaz Ali and Samuel D Silvey. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society: Series B (Methodological)*, 28(1):131–142, 1966.

4. In probability theory, an event happens almost surely if it happens with probability one.

5. $\mathbb{1}(x)$ denotes the indicator function, which returns one when x is true and zero, otherwise.

- Linas Baltrunas, Tadas Makcinskas, and Francesco Ricci. Group recommendations with rank aggregation and collaborative filtering. In ACM conference on Recommender systems, pages 119–126. ACM, 2010.
- Aharon Ben-Tal, Dick Den Hertog, Anja De Waegenare, Bertrand Melenberg, and Gijs Rennen. Robust solutions of optimization problems affected by uncertain probabilities. Management Science, 59(2):341–357, 2013.
- Anirban Bhattacharya, Debdeep Pati, Yun Yang, et al. Bayesian fractional posteriors. The Annals of Statistics, 47(1):39–66, 2019.
- Patrick Billingsley. Convergence of Probability Measures, 2nd Edition. John Wiley & Sons, 1999.
- Christopher M Bishop. Pattern Recognition and Machine Learning. Springer, 2006.
- Jose Blanchet, Yang Kang, and Karthyek Murthy. Robust wasserstein profile inference and applications to machine learning. arXiv preprint arXiv:1610.05627, 2016.
- Stephen Boyd and Lieven Vandenberghe. Convex Optimization. Cambridge university press, 2004.
- Ralph Allan Bradley and Milton E Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. Biometrika, 39(3/4):324–345, 1952.
- Francois Caron and Arnaud Doucet. Efficient Bayesian inference for generalized Bradley–Terry models. Journal of Computational and Graphical Statistics, 21(1):174–196, 2012.
- Sujoy Chatterjee, Anirban Mukhopadhyay, and Malay Bhattacharyya. A weighted rank aggregation approach towards crowd opinion analysis. Knowledge-Based Systems, 149:47–60, 2018.
- Ruidi Chen and Ioannis Paschalidis. Outlier detection using robust optimization with uncertainty sets constructed from risk measures. ACM SIGMETRICS Performance Evaluation Review, 45(3):174–179, 2018a.
- Ruidi Chen and Ioannis Ch Paschalidis. A robust learning approach for regression models based on distributionally robust optimization. The Journal of Machine Learning Research, 19(1):517–564, 2018b.
- Xi Chen, Paul N Bennett, Kevyn Collins-Thompson, and Eric Horvitz. Pairwise ranking aggregation in a crowdsourced setting. In International Conference on Web Search and Data Mining, pages 193–202. ACM, 2013.
- Kai-Yang Chiang, Cho-Jui Hsieh, and Inderjit Dhillon. Rank aggregation and prediction with item features. In Artificial Intelligence and Statistics, pages 748–756, 2017.
- Jean Daunizeau. Semi-analytical approximations to statistical moments of sigmoid and softmax mappings of normal variables. arXiv preprint arXiv:1703.00091, 2017.
- Persi Diaconis. Group representations in probability and statistics. Lecture Notes–Monograph Series, Hayward, CA: Institute of Mathematical Statistics, 11:1–192, 1988.

- Cynthia Dwork, Ravi Kumar, Moni Naor, and Dandapani Sivakumar. Rank aggregation methods for the web. In International Conference on World Wide Web, pages 613–622. ACM, 2001.
- Mohsen Ahmadi Fahandar, Eyke Hüllermeier, and Inés Couso. Statistical inference for incomplete ranking data: the case of rank-dependent coarsening. In International Conference on Machine Learning, pages 1078–1087. JMLR. org, 2017.
- Rui Gao, Xi Chen, and Anton J Kleywegt. Wasserstein distributional robustness and regularization in statistical learning. arXiv preprint arXiv:1712.06050, 2017.
- Isobel Claire Gormley and Thomas Brendan Murphy. Exploring heterogeneity in irish voting data: A mixture modelling approach. Technical report, Technical Report 05/09, Department of Statistics, Trinity College Dublin, 2005.
- Bo Han, Yuangang Pan, and Ivor W Tsang. Robust plackett–luce model for k-ary crowdsourced preferences. Machine Learning, 107(4):675–702, 2018.
- Toshihiro Kamishima. Nantonac collaborative filtering: recommendation based on order responses. In SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 583–588. ACM, 2003.
- Maurice G Kendall. A new measure of rank correlation. Biometrika, 30(1/2):81–93, 1938.
- Ashish Khetan and Sewoong Oh. Data-driven rank breaking for efficient rank aggregation. The Journal of Machine Learning Research, 17(1):6668–6721, 2016.
- Minji Kim, Farzad Farnoud, and Olgica Milenkovic. Hydra: gene prioritization via hybrid distance-score rank aggregation. Bioinformatics, 31(7):1034–1043, 2014.
- Raivo Kolde, Sven Laur, Priit Adler, and Jaak Vilo. Robust rank aggregation for gene list integration and meta-analysis. Bioinformatics, 28(4):573–580, 2012.
- Solomon Kullback and Richard A Leibler. On information and sufficiency. The Annals of Mathematical Statistics, 22(1):79–86, 1951.
- Shuai Li, Tor Lattimore, and Csaba Szepesvari. Online learning to rank with features. In International Conference on Machine Learning, pages 3856–3865, 2019.
- Xue Li, Xinlei Wang, and Guanghua Xiao. A comparative study of rank aggregation methods for partial and top ranked lists in genomic applications. Briefings in bioinformatics, 20(1):178–189, 2017.
- Yao Li, Minhao Cheng, Kevin Fujii, Fushing Hsieh, and Cho-Jui Hsieh. Learning from group comparisons: exploiting higher order interactions. In Advances in Neural Information Processing Systems, pages 4981–4990, 2018.
- Lucy Liang and Kristen Grauman. Beyond comparing image pairs: Setwise active learning for relative attributes. In IEEE conference on Computer Vision and Pattern Recognition, pages 208–215, 2014.

- Shili Lin. Rank aggregation methods. Wiley Interdisciplinary Reviews: Computational Statistics, 2 (5):555–570, 2010.
- Ao Liu, Zhibing Zhao, Chao Liao, Pinyan Lu, and Lirong Xia. Learning Plackett-Luce mixtures from partial preferences. In AAAI Conference on Artificial Intelligence, volume 33, pages 4328–4335, 2019.
- RD Luce. Individual choice theory: A theoretical analysis, 1959.
- Colin Mallows. Non-null ranking models. Biometrika, 1957.
- Jeffrey W Miller and David B Dunson. Robust Bayesian inference via coarsening. Journal of the American Statistical Association, 114(527):1113–1125, 2019.
- Cristina Mollica and Luca Tardella. Bayesian plackett–luce mixture models for partially ranked data. psychometrika, 82(2):442–458, 2017.
- Hongseok Namkoong and John C Duchi. Variance-based regularization with convex objectives. In Advances in Neural Information Processing Systems, pages 2971–2980, 2017.
- Yuangang Pan, Bo Han, and Ivor W. Tsang. Stagewise learning for noisy k-ary preferences. Machine Learning, 107(8-10):1333–1361, 2018.
- Yuangang Pan, Ivor W. Tsang, Avinash K. Singh, Chin-Teng Lin, and Masashi Sugiyama. Stochastic multichannel ranking with brain dynamics preferences. Neural Computation, 32(8):1499–1530, 08 2020. ISSN 0899-7667.
- Yuangang Pan, Ivor W. Tsang, Yueming Lyu, Avinash K. Singh, and Chin-Teng Lin. Online mental fatigue monitoring via indirect brain dynamics evaluation. Neural Computation, 33(6):1616–1655, 05 2021. ISSN 0899-7667.
- Robin L Plackett. The analysis of permutations. Journal of the Royal Statistical Society: Series C (Applied Statistics), 24(2):193–202, 1975.
- Karthik Raman and Thorsten Joachims. Methods for ordinal peer grading. In SIGKDD International Conference on Knowledge Discovery and Data Mining, pages 1037–1046. ACM, 2014.
- Philip H Ramsey. Critical values for spearman’s rank order correlation. Journal of educational statistics, 14(3):245–253, 1989.
- Vikas C Raykar, Shipeng Yu, Linda H Zhao, Gerardo Hermosillo Valadez, Charles Florin, Luca Bogoni, and Linda Moy. Learning from crowds. Journal of Machine Learning Research, 11(Apr): 1297–1322, 2010.
- Mehdi SM Sajjadi, Morteza Alamgir, and Ulrike von Luxburg. Peer grading in a course on algorithms and data structures: Machine learning algorithms do not improve over simple baselines. In ACM Conference on Learning@ Scale, pages 369–378. ACM, 2016.
- Sundararajan Sellamanickam, Priyanka Garg, and Sathiya Keerthi Selvaraj. A pairwise ranking based approach to learning with positive and unlabeled examples. In International Conference on Information and Knowledge Management, pages 663–672. ACM, 2011.

- Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. arXiv preprint arXiv:1710.10571, 2017.
- David Spiegelhalter, Nicky Best, and Bradley Carlin. Bayesian deviance, the effective number of parameters, and the comparison of arbitrarily complex models. Journal of Royal Statistical Society, 64, 05 1998.
- Louis L Thurstone. A law of comparative judgment. Psychological review, 34(4):273, 1927a.
- Louis L Thurstone. The method of paired comparisons for social values. The Journal of Abnormal and Social Psychology, 21(4):384, 1927b.
- C. Villani. Optimal Transport: Old and New. Grundlehren der mathematischen Wissenschaften. Springer Berlin Heidelberg, 2008. ISBN 9783540710509.
- Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In Advances in Neural Information Processing Systems, pages 5334–5344, 2018.
- Ruby C Weng and Chih-Jen Lin. A Bayesian approximation method for online ranking. Journal of Machine Learning Research, 12(Jan):267–300, 2011.
- Qianqian Xu, Ming Yan, Chendi Huang, Jiechao Xiong, Qingming Huang, and Yuan Yao. Exploring outliers in crowdsourced ranking for qoe. In International Conference on Multimedia, pages 1540–1548. ACM, 2017.
- Zhibing Zhao and Lirong Xia. Learning mixtures of plackett-luce models from structured partial orders. In Advances in Neural Information Processing Systems, pages 10143–10153, 2019.
- Zhibing Zhao, Peter Piech, and Lirong Xia. Learning mixtures of Plackett-Luce models. In International Conference on Machine Learning, pages 2906–2914. PMLR, 2016.