



**HAL**  
open science

# Transductive Conformal Inference for Ranking

Jean-Baptiste Fermanian, Pierre Humbert, Gilles Blanchard

► **To cite this version:**

Jean-Baptiste Fermanian, Pierre Humbert, Gilles Blanchard. Transductive Conformal Inference for Ranking. 2025. hal-04895043

**HAL Id: hal-04895043**

**<https://hal.science/hal-04895043v1>**

Preprint submitted on 17 Jan 2025

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Transductive Conformal Inference for Ranking

Jean-Baptiste Fermanian\*   Pierre Humbert†   Gilles Blanchard‡

## Abstract

We introduce a method based on Conformal Prediction (CP) to quantify the uncertainty of full ranking algorithms. We focus on a specific scenario where  $n + m$  items are to be ranked by some “black box” algorithm. It is assumed that the relative (ground truth) ranking of  $n$  of them is known. The objective is then to quantify the error made by the algorithm on the ranks of the  $m$  new items among the total  $(n + m)$ . In such a setting, the true ranks of the  $n$  original items in the total  $(n + m)$  depend on the (unknown) true ranks of the  $m$  new ones. Consequently, we have no direct access to a calibration set to apply a classical CP method. To address this challenge, we propose to construct distribution-free bounds of the unknown conformity scores using recent results on the distribution of conformal p-values. Using these scores upper bounds, we provide valid prediction sets for the rank of any item. We also control the false coverage proportion, a crucial quantity when dealing with multiple prediction sets. Finally, we empirically show on both synthetic and real data the efficiency of our CP method.

## 1 Introduction

Ranking is a fundamental problem in machine learning where the objective is to sort some items, such as documents or products, by their relevance to a query or user profile. It has a wide range of applications, ranging from document retrieval (Cao, Xu, et al., 2006), collaborative filtering (Liu and Yang, 2008), sentiment analysis (Liu, Bi, et al., 2017), and product rating. While there exist a lot of ranking algorithms in the literature (see e.g. Liu et al. (2009) for a review), the quantification of the uncertainty in their rank predictions using conformal prediction is relatively new (Angelopoulos, Krauth, et al., 2023; Xu et al., 2024). Indeed, so far, the design of ranking algorithms has mainly focused on the training phase and their evaluations regarding the uncertainty

---

\*Université de Montpellier, Inria, France.

†Sorbonne Université and Université Paris Cité, CNRS, Laboratoire de Probabilités, Statistique et Modélisation

‡Université Paris Saclay, Institut Mathématique d’Orsay, France.

of their predicted ranking, is not often studied. With the increasing popularity of ranking methods (especially in the learning-to-rank literature Liu et al., 2009), this evaluation is however a mandatory step to deploy them for real-world applications.

Our main objective is to quantify the uncertainty of ranking algorithms by constructing distribution free prediction set for the ranks of each items. Formally, given  $n + m$  items with an underlying unknown total order, we want to construct prediction sets that contains the true rank of each item with high probability. We consider a particular scenario where the first  $n$  items have known relative rankings, but their positions within the whole  $n + m$  items are unknown. Furthermore, the remaining  $m$  items have no known ranking information.

This scenario occurs for example in a matchmaking problem (Alman and McKay, 2017; Herbrich, Minka, et al., 2006; Minka et al., 2018). In that framework, we observe the outcome of matches between  $n$  players and then deduce a ranking of these players. When  $m$  new players enter the game, we want to predict their rank among all previous players in order to organize matches between players with the same skills. Another example is when we want to replace an expensive ranking process with a cheaper but less reliable one (based, for example, on a smaller number of item comparisons). In this situation, we want to estimate the reliability of the cheaper process compared to the costly one considered as a proxy truth. We can therefore randomly select  $n$  items, consider that their relative true ranks are the rank returned by the more efficient process and use them to quantify the uncertainty of the other one. As a final motivating example, consider a problem where we want to predict which types of  $n + m$  individuals will best adapt to a given ecosystem. Since environmental variations can affect the whole population, instead of predicting an absolute measure of adaptation (e.g. height, weight, health indicators), we are more interested in a relative ranking of these individuals. To this end, a scientist can sample a subset of  $n$  individuals to measure their characteristics and obtain their relative true ranks. Then, the scientist can use a ranking algorithm to rank all the individuals and use the  $n$  first points to quantify the uncertainty in her ranking.

Formally, for each item  $i \in \llbracket n + m \rrbracket$ , we want to construct a *marginally valid* set which contains, with high probability, its unknown rank  $R_i$  inside the  $n + m$  items:

$$\mathbb{P}\left(R_i \in \hat{\mathcal{C}}_i\right) \geq 1 - \alpha, \tag{1}$$

where  $\alpha \in (0, 1)$  is a desired miscoverage level. In addition, as we construct prediction sets for multiple points, we also want to control the false coverage proportion, which is the average number of mispredictions. More precisely, we want to construct  $m$

prediction sets  $(\widehat{\mathcal{C}}_i)_{i=1}^m$  satisfying for some small  $\beta > 0$ :

$$\mathbb{P}\left(\frac{1}{m} \sum_{i=1}^m \mathbf{1}\{R_i \notin \widehat{\mathcal{C}}_i\} \leq \alpha\right) \geq 1 - \beta. \quad (2)$$

There exist several methods to construct marginally valid sets in regression or classification settings, based on the so-called Conformal Prediction (CP) framework (Papadopoulos et al., 2002; Romano et al., 2019; Vovk et al., 2005), but these do not directly apply in this specific setting of ranking. Indeed, the problem of ranking can be seen as a regression problem where the outcome, the rank, depends on all the points observed, or as a classification problem with a moving number of classes. An important difference with a classical CP framework is that we have not access directly to a calibration set, as knowing the relative rank of the  $n$  first items is not sufficient to know their ranks within the whole sample.

**Contributions.** This work constructs marginally valid prediction sets, satisfying Eq. (1), for the ranks of  $n + m$  items. We assume already having a ranking algorithm  $\mathcal{A}$  (considered as a “black box”) and knowing the exact relative ranks of  $n$  items. These  $n$  items will be used to quantify the uncertainty of the algorithm  $\mathcal{A}$ . Our main contribution is to propose a general method for this calibration phase. The difficulty is that we cannot directly quantify the algorithm’s error on the calibration set, as the true ranks are only known through the relative ranks. We then provide computable bounds of these ranks, supported by theoretical arguments, to construct valid prediction sets for which we also control the false coverage proportion. These findings provide a practical solution to the challenges of constructing valid prediction sets in ranking problems, where traditional CP methods face limitations due to the interdependence of ranks.

## 2 Background

### 2.1 Split conformal prediction

Conformal Prediction (CP) is a framework introduced by Vovk et al. (2005) to construct distribution-free prediction sets quantifying the uncertainty in the predictions of an algorithm. Formally, given a calibration set  $\{(X_i, Y_i)\}_{i \in \llbracket n \rrbracket}$  of points in  $\mathcal{X} \times \mathcal{Y}$  and an algorithm  $\mathcal{A}$ , a CP method constructs for a new point  $X_{n+1} \in \mathcal{X}$  a set  $\widehat{\mathcal{C}}$  containing the unobserved outcome  $Y_{n+1} \in \mathcal{Y}$  with high probability:

$$\mathbb{P}\left(Y_{n+1} \in \widehat{\mathcal{C}}(X_{n+1})\right) \geq 1 - \alpha. \quad (3)$$

In classification,  $\mathcal{Y}$  is the set of classes and  $\mathcal{A}$  a classifier and in regression,  $\mathcal{Y} = \mathbb{R}$  and  $\mathcal{A}$  is a regressor. The idea behind CP methods is to first build non-conformity scores  $V_i := s(X_i, Y_i)$  on the calibration set, where  $s : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}$ , is a score function which is intended to quantify the error of the algorithm at a given point. Many score functions have been considered to catch different type of information, see e.g. Angelopoulos, Bates, et al. (2023) for a review. In regression, a common choice is to use the absolute residual  $s(x, y) = |y - \mathcal{A}(x)|$  where the algorithm  $\mathcal{A}$  can either be trained on an independent sample (split conformal prediction, Papadopoulos et al., 2002) or retrained on a subset of points (full conformal prediction Vovk et al., 2005). In this paper, we will focus on the case where an already trained algorithm is provided, which corresponds to the case of split CP. Then, for this score function and some integer  $k$ , the prediction set is

$$\widehat{\mathcal{C}}_k(x) := \left\{ y \in \mathcal{Y} : s(x, y) \leq V_{(k)} \right\}, \quad (4)$$

where  $V_{(k)}$  is the  $k$ -th smallest score of  $V_1, \dots, V_n$ . This set is marginally valid under mild assumptions as presented in the following theorem.

**Theorem 1** (Lei et al., 2018; Vovk et al., 2005). *If the scores  $V_1, \dots, V_{n+1}$  are exchangeable, then for any  $\alpha \in (0, 1)$  and  $k = \lceil (1 - \alpha)(n + 1) \rceil$  the set (4) returned by the split CP method satisfies:*

$$\mathbb{P}\left(Y_{n+1} \in \widehat{\mathcal{C}}_k(X_{n+1})\right) \geq 1 - \alpha.$$

We refer to Vovk et al. (2005), Angelopoulos, Bates, et al. (2023) and Fontana et al. (2023) for in-depth presentations of CP. See also Manokhin (2024) for a curated list of CP papers.

## 2.2 False coverage proportion

Some recent results have considered the transductive setting (Vovk, 2013) where conformal prediction sets are constructed for  $m \geq 2$  test points  $\{(X_{n+i}, Y_{n+i})\}_{i=1}^m$ . In this case, while maintaining the coverage guarantee (3), an usual goal is to control the False Coverage Proportion (FCP) defined as:

$$\text{FCP} := \frac{1}{m} \sum_{i=1}^m \mathbf{1}\left\{Y_{n+i} \notin \widehat{\mathcal{C}}(X_{n+i})\right\}, \quad (5)$$

where  $\widehat{\mathcal{C}}(X_{n+i})$  are some prediction sets. They can be constructed as in Eq. (4) for instance. In words, the FCP is the proportion of observations in the test set that fall outside the constructed prediction sets.

If the prediction sets are of the form (4), under the assumptions of Theorem 1 we have  $\mathbb{E}[\text{FCP}] \leq \alpha$ . Furthermore, the exact distribution of the FCP is known (Huang et al., 2024; Marques F., 2024). Going further, Gazin et al. (2024) studied the full joint distribution of the insertion ranks of test scores into calibration scores, leading under the same assumptions to a uniform control of the FCP over the parameter  $k$  in (4), in probability. The asymptotic behaviour of the FCP in the form of a functional CLT has also been studied recently by Gazin (2024). With all these theoretical results, it is possible to control in probability the FCP by adjusting the coverage of the prediction sets.

### 2.3 Related work in ranking

Given a set of items the goal of ranking is to infer a particular ordering over these items.

There exist a multitude of settings and approaches to solve the ranking problem. In the learning-to-rank literature, and more specifically, in the pairwise approach, the ranking task is formalised as a classification of pairs of items into two classes: correctly ranked and incorrectly ranked. Herbrich, Graepel, et al. (1999) proposed an algorithm based on Support Vector Machine (SVM) called RankSVM to solve this classification problem. Alternative techniques can be used such as boosting with RankBoost (Freund et al., 2003), or Neural Network with RankNet (Burges et al., 2005). Other methods based on the pointwise or listwise approaches are also widely study (Cao, Qin, et al., 2007; Li, 2011; Liu et al., 2009). Another line of research solve the ranking problem through binary comparisons of items. The different settings considered allow these binary comparisons to be measured either completely at random (Wauthier et al., 2013), actively (Ailon, 2012; Heckel et al., 2019), or several times (Feige et al., 1994; Negahban et al., 2012). They can also assume that all the possible comparisons between the items are known but up to some noise (Braverman and Mossel, 2007, 2009).

Overall, the quantification of uncertainty of these ranking algorithms has been little studied in CP. Recently, Angelopoulos, Krauth, et al., 2023 have developed a CP method for quantifying the uncertainty of learning-to-rank algorithms specially tailored for recommendation systems. Xu et al., 2024 also proposed to use CP, but for the ranked retrieval task. While related to the present work, they do not consider the same setting.

### 3 Conformal ranking

We now describe our setting, as well as our methodology and main results to provide valid prediction sets.

**Notations.** For a point  $y$  in  $\mathbb{R}$  and a finite subset  $\mathcal{D} \subset \mathbb{R}$ , the rank of  $y$  in  $\mathcal{D}$  is:

$$\mathbf{R}(y, \mathcal{D}) := \sum_{z \in \mathcal{D}} \mathbf{1}\{y \geq z\}.$$

For an index  $1 \leq i \leq |\mathcal{D}|$ , we denote the point of  $\mathcal{D}$  of rank  $i$  by:

$$\mathbf{R}^{-1}(i, \mathcal{D}) := \sum_{z \in \mathcal{D}} z \mathbf{1}\{\mathbf{R}(z, \mathcal{D}) = i\}.$$

A multiset (or bag) will be denoted  $\wr \cdot \wr$  and is an unordered set with potentially repeated elements.

#### 3.1 Setting

Let us consider  $n + m$  items, each associated with a pair  $(X_i, Y_i) \in \mathcal{X} \times \mathbb{R}$ . We suppose that these items are divided into two sets: a calibration set of size  $n$ ,  $\{(X_i, Y_i)\}_{i \in \llbracket n \rrbracket}$  and a test set of size  $m$ ,  $\{(X_{n+i}, Y_{n+i})\}_{i \in \llbracket m \rrbracket}$ . For any item  $i$ ,  $X_i$  is, for instance, a vector of features, while  $Y_i$  allows us to properly define its true rank within each subset or the entire set through the following variables:

$$\begin{aligned} R_i^c &:= \mathbf{R}(Y_i, \wr Y_j \wr_{j \in \llbracket n \rrbracket}), & R_i^t &:= \mathbf{R}(Y_i, \wr Y_{n+j} \wr_{j \in \llbracket m \rrbracket}), \\ R_i^{c+t} &:= \mathbf{R}(Y_i, \wr Y_j \wr_{j \in \llbracket n+m \rrbracket}) = R_i^c + R_i^t. \end{aligned} \tag{6}$$

We assume throughout the paper that:

**Assumption 1.** The vector  $(Y_i)_{i \in \llbracket n+m \rrbracket}$  is *exchangeable* and has *no ties*. In particular, this implies that there is a *total order* between the  $n + m$  items.

The values  $\{Y_i\}_i$  can be interpreted as an *unobservable underlying truth*. In our setting, we insist that only the features  $(X_i)_{i \in \llbracket n+m \rrbracket}$  and the true ranks of the calibration points  $\{R_i^c\}_{i \in \llbracket n \rrbracket}$  are observed. For example, they can model the intrinsic skill of some players only observable by comparing them with each other. Assumption 1 is needed to model correctly our problem of full ranking. The following assumption is essential to be able to construct conformal prediction sets for the ranks.

**Assumption 2.** The vector  $(X_i, R_i)_{1 \leq i \leq \llbracket n+m \rrbracket}$  is exchangeable.

To estimate the true ranks  $R_1^{c+t}, \dots, R_{n+m}^{c+t}$ , we suppose having an algorithm  $\mathcal{A} : \mathcal{X} \times \mathcal{X}^{n+m} \rightarrow \mathbb{K}$  which is intended to predict the rank of a point inside a set of points. Its inputs are the target item and the multiset of items among which it seeks to sort it. For clarity, when there is no confusion, we will sometimes omit the dependence in the multiset of items:  $\mathcal{A}(x) := \mathcal{A}(x, \wr x_\ell \wr)$ . This algorithm can predict for each point either a *rank* or a *value* from which a rank can be deduced. For these two situations, we propose two conformity score functions which quantify the error made by  $\mathcal{A}$  in its ranking:

**(RA) setting:**  $\mathbb{K} = \mathbb{N}$ , the rank is directly predicted by the algorithm  $\mathcal{A}$ :  $\widehat{R}_i^{c+t} := \mathcal{A}(X_i, \wr X_j \wr_{j \in \llbracket n+m \rrbracket})$ .

In this situation, we can consider the classical residual scores, i.e., for  $x, (x_j)_j$  in  $\mathcal{X}$  and  $r \in \mathbb{N}$ :

$$s^{\mathbf{RA}}(x, r, \wr x_j \wr_j) := |r - \mathcal{A}(x, \wr x_j \wr_j)|.$$

This is simply the absolute difference between  $r$  and the predicted rank of  $x$  inside  $(x_j)_j$ .

**(VA) setting:**  $\mathbb{K} = \mathbb{R}$ , the predicted rank is deduced from values constructed by the algorithm  $\mathcal{A}$ :

$$\widehat{R}_i^{c+t} := \mathbf{R} \left( \mathcal{A}(X_i, \wr X_\ell \wr_\ell), \left\{ \mathcal{A}(X_j, \wr X_\ell \wr_\ell) \right\}_{j \in \llbracket n+m \rrbracket} \right).$$

Here, because  $\mathcal{A}(x)$  is a real value, we consider a more refined score function than  $s^{\mathbf{RA}}$ , for  $x, (x_j)_j$  in  $\mathcal{X}$  and  $r \in \mathbb{N}^*$ :

$$s^{\mathbf{VA}}(x, r, \wr x_j \wr_j) := \left| \mathbf{R}^{-1} \left( r, \wr \mathcal{A}(x_j) \wr_j \right) - \mathcal{A}(x) \right|.$$

This score quantifies the distance between the value attributed by the algorithm to the point  $x$  and the value it should have had to be at rank  $r$ . We restrict ourselves to these two scores for our experiments, but our methodology and theoretical results are more generally valid for any score function  $s(x, r, \wr x_j \wr_j)$ .

*Remark 2.* The fact that the algorithm depends on the multiset rather than an  $(n+m)$ -tuple is important, as it implies that the algorithm treats the points exchangeably. With Assumption 2, the vector  $(\mathcal{A}(X_i))_{i \in \llbracket n+m \rrbracket}$  is thus exchangeable.

*Remark 3.* By considering the predicted ranks as scores, we see that **(RA)** is a sub-case of **(VA)**. Moreover, as in **(RA)**,  $\mathbf{R}^{-1}(r, \wr \mathcal{A}(x_i) \wr_i)_i = r$ , the function  $s^{\mathbf{VA}}$  reduces to the residual function  $s^{\mathbf{RA}}$ .



### 3.2 Methodology and main results

We now introduce our main theoretical results to construct prediction sets for the ranks  $R_1^{c+t}, \dots, R_{n+m}^{c+t}$ . Recall that our objective is to build for  $j \in \llbracket n+m \rrbracket$  marginally sets  $\widehat{\mathcal{C}}_j$  for  $R_j^{c+t}$ , i.e. satisfying Eq. (1) for a given  $\alpha \in (0, 1)$ . As we construct multiple prediction sets, we also provide a control of the FCP (Eq.(5)).

To construct these sets, a natural strategy is to calculate the quantile of order  $1 - \alpha$  of the scores  $s(X_i, R_i^{c+t}, \wr \mathcal{A}(X_j)_{j \in \llbracket n+m \rrbracket})$  computed on the calibration set as this is done in the split CP method (see Section 2.1). However, these scores depend on the quantities  $\{R_i^{c+t}\}_{i \in \llbracket n \rrbracket}$  which are unknown. Consequently, they are not computable and we must find another strategy. The following result shows that, if we can bound the ranks of the calibration  $(R_i^{c+t})_{i \in \llbracket n \rrbracket}$  with high probability, then it is possible to construct marginally valid sets for the test ranks  $\{R_{n+i}^{c+t}\}_{i \in \llbracket m \rrbracket}$ .

**Theorem 4.** *Let  $\alpha \in (0, 1)$ ,  $\delta < \alpha$  and  $R_i^-, R_i^+ \in \mathbb{R}$  ( $i \in \llbracket n \rrbracket$ ) be random variables depending on  $(X_i, R_i^c)_{i \in \llbracket n \rrbracket}$  and satisfying*

$$\mathbb{P}(\forall i \in \llbracket n \rrbracket : R_i^{c+t} \in [R_i^-, R_i^+]) \geq 1 - \delta. \quad (7)$$

Let us define the proxy scores for  $i \in \llbracket n \rrbracket$ :

$$S_i := \max_{r \in [R_i^-, R_i^+]} s(X_i, r, \wr X_j)_{j \in \llbracket n+m \rrbracket} \quad (8)$$

and the associated conformal prediction set:

$$\widehat{\mathcal{C}}_k(x) = \{r : s(x, r, \wr X_i)_{i \in \llbracket n+m \rrbracket}\} \leq S_{(k)} \}, \quad (9)$$

$S_{(k)}$  being the  $k$ -th smallest proxy score in  $S_1, \dots, S_n$ .

If Assumption 2 is satisfied, then for  $k = \lceil (1 - \alpha + \delta)(n + 1) \rceil$  and any  $i \in \llbracket m \rrbracket$ :

$$\mathbb{P}(R_{n+i}^{c+t} \in \widehat{\mathcal{C}}_k(X_i)) \geq 1 - \alpha.$$

For the score functions  $s^{\mathbf{RA}}$  and  $s^{\mathbf{VA}}$ , the proxy scores of Eq. (8) and the associated conformal sets have explicit expressions which make them easy to compute.

**Example 5** ( $s = s^{\mathbf{RA}}$ ). The maximum (8) is necessarily attained at the edges of the interval  $[R^-, R^+]$ :

$$\max_{r \in [R^-, R^+]} s^{\mathbf{RA}}(x, r, \wr x_i) = \max(|R^- - \mathcal{A}(x, \wr x_i)|, |R^+ - \mathcal{A}(x, \wr x_i)|).$$

The associated conformal set is:

$$\widehat{\mathcal{C}}^{\mathbf{RA}}(x) = [\mathcal{A}(x, \wr X_i) \pm S_{(k)}].$$

**Example 6** ( $s = s^{\text{VA}}$ ). Similarly as the previous example, the maximum is attained at the edges of  $[R^-, R^+]$ :

$$\max_{r \in [R^-, R^+]} s^{\text{VA}}(x, r, \wr_{x_i} \wr_i) = \max_{r \in \{R^-, R^+\}} (|\mathcal{A}(x) - \mathbf{R}^{-1}(r, \wr_{\mathcal{A}(x_i)} \wr_i)|).$$

The corresponding conformal set is then the interval

$$\widehat{\mathcal{C}}^{\text{VA}}(x) = \left[ \mathbf{R} \left( \mathcal{A}(x) - S_{(k)}, \wr_{\mathcal{A}(x_i)} \wr_i \right), \mathbf{R} \left( \mathcal{A}(x) + S_{(k)}, \wr_{\mathcal{A}(x_i)} \wr_i \right) \right].$$

We see that contrary to the conformal intervals  $\widehat{\mathcal{C}}^{\text{RA}}$  which have the same length for any point  $x$ , the length of the conformal interval  $\widehat{\mathcal{C}}^{\text{VA}}$  can vary depending on the amount of points around  $\mathcal{A}(x)$ .

The following result gives a control in high probability of the false coverage rate on the test sample.

**Proposition 7.** *Let  $\alpha \in (0, 1)$  and consider the prediction sets  $\widehat{\mathcal{C}}_k$  defined in Eq. (9). Assume Assumption 2 holds, let  $k = \lceil (1 - \alpha)(n + 1) \rceil$ , then with probability at least  $1 - \beta - \delta$ :*

$$\frac{1}{m} \sum_{i=1}^m \mathbf{1} \left\{ R_{n+i}^{c+t} \notin \widehat{\mathcal{C}}_k(X_{n+i}) \right\} \leq \alpha + \lambda_{n,m}, \quad (10)$$

where  $\lambda_{n,m} = \sqrt{\frac{\log(C\sqrt{\tau_{n,m}}/\beta)}{\tau_{n,m}}}$ ,  $\tau_{n,m} = \frac{nm}{n+m}$  and  $C = 4\sqrt{2\pi}$  works.

When the sample sizes increase, then  $\lambda_{n,m} \rightarrow 0$  and the average error on the test sample is at most  $\alpha$ . In practice, using  $\alpha' = \alpha - \lambda_{n,m}$  to control the FCP at level  $\alpha$  is too conservative as the constant  $C$  is not sharp enough. We instead use the numerical procedure explained by Gazin et al. (2024, Remark 2.6) to find a sharper but implicit value for  $\lambda_{n,m}$  (see Appendix B for details).

*Remark 8.* We have focused in this section on the construction and the guarantees of prediction sets for the ranks of the test points. For the calibration points, we can directly use the high probability bounds (7) on which our method is built. Indeed, by construction the sets  $[R_i^-, R_i^+]$  have a null FCP with probability at least  $1 - \delta$  for  $i \in \llbracket n \rrbracket$ . Their explicit construction is described in the following section.

## 4 Rank of the calibration points

To make our method complete, it remains to find  $R_i^-$  and  $R_i^+$  satisfying Eq. (7). Let us recall that  $R_i^{c+t} = R_i^c + R_i^t$ . Hence, we know that  $R_i^{c+t} \in [R_i^c, R_i^c + m]$  and we can use Theorem 4 with  $R_i^- = R_i^c$  and  $R_i^+ = R_i^c + m$ . However, using these bounds for instance with the residual conformity function  $s^{\text{RA}}$ , will increase the length of the conformal set  $\widehat{\mathcal{C}}^{\text{RA}}(x)$  by  $m$ . In this section, we show that it is possible to find theoretical bounds and practical bounds satisfying Eq. (7) that do not excessively increase the length of the final conformal sets.

### 4.1 Theoretical bound

As the calibration and test values  $(Y_i)_{i \in \llbracket n+m \rrbracket}$  are exchangeable (Assumption 1), we can expect that a calibration point with a low rank in the calibration set will have a low rank in the test set. The following proposition justifies this intuition.

**Proposition 9.** *Under Assumption 1, for  $\delta \in (0, 1)$ :*

$$\mathbb{P}\left(R_i^{c+t} \in [R_i^-, R_i^+] \cap \llbracket n+m \rrbracket, i \in \llbracket n \rrbracket\right) \geq 1 - \delta,$$

where

$$R_i^\pm = R_i^c + (m+1) \left( \frac{R_i^c}{n} \pm \sqrt{\frac{\log(C \sqrt{\tau_{n,m}} / \delta)}{\tau_{n,m}}} \right),$$

$\tau_{n,m} = \frac{nm}{n+m}$ , and  $C = 4\sqrt{2\pi}$  works.

The proof can be found in the supplementary and relies on Gazin et al. (2024, Theorem 2.4). Relatively to the naive idea where we get an interval of length  $m$  for the rank  $R^{c+t}$ , we obtain here an interval of length  $O(m/\sqrt{\tau_{n,m}})$ . The length have been thus reduced by a factor  $\sqrt{\tau_{n,m}} \simeq \min(\sqrt{n}, \sqrt{m})$ .

### 4.2 Numerical bounds

The explicit bounds given by Proposition 9 are often too conservative in practice. In this section, we construct sharp numerical bounds  $R_i^-, R_i^+$  for the ranks  $R_i^{c+t}$  such that the event

$$\{\forall 1 \leq j \leq n, R_j^{c+t} \text{ is between } R_i^- \text{ and } R_i^+\}$$

holds with high probability. Let us denote for  $r \in \llbracket n \rrbracket$ ,  $R_{(r)}^{c+t}$  the rank in the whole sample of the item of rank  $r$  in the calibration. In particular,  $R_{(R_i^c)}^{c+t} = R_i^{c+t}$ . Our

idea is to use that the random vector  $R_{srt}^{c+t} = (R_{(1)}^{c+t}, \dots, R_{(n)}^{c+t})$  follows an universal distribution independent of the distribution of the data to simulate them and to estimate an envelope. This distribution can be drawn using Algorithm 1 thanks to the exchangeability of the data (Assumption 1; see Gazin et al., 2024 for more details). An envelope can then be estimated by a Monte-Carlo approach using the empirical cumulative distribution function.

---

**Algorithm 1** Simulation of  $R_{srt}^{c+t}$

---

- 1: Draw  $n + m$  uniform random variable  $U_i$  on  $[0, 1]$ .
  - 2: **for**  $i = 1, \dots, n$  **do**
  - 3:      $\tilde{R}_i^{c+t} \leftarrow \mathbf{R}(U_i, \{U_j\}_{j \in \llbracket n+m \rrbracket})$
  - 4: **end for**
  - 5: **Output:** Sort  $\tilde{R}^{c+t}$
- 

We consider two forms of envelope in this work, the linear and the quantile detailed below, which are parameterized by real factors  $c$  and  $\gamma$ . To estimate these envelopes, we draw  $K$  realizations of vectors following the same distribution as  $(R_{srt}^{c+t})$  by repeating  $K$  times Algorithm 1 and estimate the (multidimensional) cumulative function of this sample. Then, we optimize the parameter  $c$  or  $\gamma$  for minimizing the width of the envelope (which depends directly of this parameter) while containing a proportion  $1 - \delta$  of the  $K$  trajectories. For the quantile envelope, this procedure is detailed in Algorithm 2.

**Linear envelope:** We consider an envelop of the form of Proposition 9:

$$\widehat{R}_i^\pm = R_i^c + (m + 1) \left( \frac{R_i^c}{n} \pm \widehat{c} \right),$$

where  $\widehat{c}$  is minimized under the condition that a proportion  $(1 - \delta)$  of the generated vectors are in the envelop. We then take the minimum with  $n + m$  and the maximum with 1 to keep the bounds in  $\llbracket n + m \rrbracket$ .

**Quantile envelope:** For an item of rank  $r = R_i^c$ , we propose to choose the bound  $\widehat{R}_i^+$  (respectively  $\widehat{R}_i^-$ ) as the empirical quantile of order  $1 - \widehat{\gamma}$  (resp.  $\widehat{\gamma}$ ) of the simulated ranks  $R_{(r)}^{c+t}$ , i.e. the ranks in the whole sample of the points of rank  $r$  in the calibration. As we want an uniform control of these bounds, we need to adjust  $\gamma$ .

Hence, the parameter  $\gamma$  is maximized under the condition that a proportion  $1 - \delta$  of the generated vectors are in the envelop. Algorithm 2 summarizes the full process to construct this envelope. In this algorithm, for  $x \in \mathbb{R}^K$  and  $u \in (0, 1)$ ,  $\widehat{Q}_u(x)$  denotes

the empirical quantile of order  $u$  of the sample  $x$ .

---

**Algorithm 2** Quantile envelope procedure
 

---

- 1: **Input:**  $n, m, K$ , and  $\delta \in (0, 1)$ ,  $(R_i^c)_{i \in [n]}$
  - 2: **for**  $k = 1, \dots, K$  **do**
  - 3:      $\tilde{R}^{(k)} \leftarrow$  Output of **Algorithm 1**
  - 4: **end for**
  - 5:  $\hat{\gamma} \leftarrow \max \gamma$  such that:  
 $\frac{1}{K} \sum_k \mathbf{1} \left\{ \exists i : \tilde{R}_i^{(k)} \notin \left[ \hat{Q}_\gamma \left( (\tilde{R}_i^{(\ell)})_{\ell \in [K]} \right), \hat{Q}_{1-\gamma} \left( (\tilde{R}_i^{(\ell)})_{\ell \in [K]} \right) \right] \right\} \leq \delta$
  - 6: **for**  $j = 1, \dots, n$  **do**
  - 7:      $\hat{R}_j^- \leftarrow \hat{Q}_{\hat{\gamma}} \left( (\tilde{R}_j^{(\ell)})_{\ell \in [K]} \right)$
  - 8:      $\hat{R}_j^+ \leftarrow \hat{Q}_{1-\hat{\gamma}} \left( (\tilde{R}_j^{(\ell)})_{\ell \in [K]} \right)$
  - 9: **end for**
  - 10: **Output:**  $\hat{R}_{R_i^c}^-, \hat{R}_{R_i^c}^+$  for  $i \in [n]$ .
- 

Figure 1 displays these two envelopes for  $n = 50$ ,  $m = 500$ , and  $\delta = 0.1$ . The blue lines are the linear envelope and the red ones are the quantile envelope. The quantile envelope is smaller for the low and high rank points while the linear envelope is smaller for the points of medium rank.

**Monte-Carlo guarantees.** The estimation of the envelopes brings an other source of error in our procedure. In the following proposition, we show that this error can be easily controlled and the parameter  $K$  can be chosen large enough to make it negligible.

**Proposition 10.** *Let  $K, n, m \in \mathbb{N}^*$  and  $(R^{(k)})_{k \in [K]}$  a  $K$ -sample of ordered ranks drawn from Algorithm 1 with parameters  $n$  and  $m$ . If  $(\hat{R}_i^\pm)_{i \in [n]}$  is an estimated envelop, possibly depending on  $(R^{(k)})_k$  such that, almost surely:*

$$\frac{1}{K} \sum_{k=1}^K \mathbf{1} \left\{ \exists i : R_i^{(k)} \notin \left[ \hat{R}_i^-, \hat{R}_i^+ \right] \right\} \leq \delta,$$

then:

$$\mathbb{P} \left[ \forall i : R_i^{c+t} \in \left[ \hat{R}_{R_i^c}^-, \hat{R}_{R_i^c}^+ \right] \right] \geq 1 - \delta - 4 \sqrt{\frac{\log nK}{K}}. \quad (11)$$

The proof is given in Appendix A.5. Notice that this result is not specific to the linear or the quantile envelope we consider but rather to the Monte-Carlo procedure. It is




Figure 1: The two envelopes for  $n = 50$ ,  $m = 500$ ,  $\delta = 0.1$  and  $K = 100000$ . The blue and red lines represent respectively the linear and quantile envelopes, and each black curve is an ordered realisation of  $R^{c+t}$ .

indeed based on a control of the multi-dimensional empirical cumulative distribution function from Naaman (2021).

## 5 Experiments

### 5.1 Synthetic data

In this section, we empirically evaluate our CP method on synthetic data. We compare its performance with the “optimal” conformal method that uses non-conformity scores calculated with the unobserved true ranks  $\{R_i^{c+t}\}_{i \in [n]}$ . This method cannot be used in real scenarios since we need to know the true ranks of the items in the calibration set, but we consider it as an oracle. These two methods are hereafter referred to as Transductive Conformal Prediction for Ranking (TCPR) and Oracle. The code of our method is available at <https://github.com/pierreHmbt/transductive-conformal-inference-for-ranking>.

**Data generation.** We use the same generation process as Pedregosa et al. (2012) with slight modifications to match our setting. In detail, we draw  $n + m$  pair of points

## Transductive Conformal Inference for Ranking




Figure 2: Synthetic data: FCP and relative lengths obtained for RankSVM and BRE-0.01 on the **(RA)** and **(VA)** settings when  $m = 500$  and  $n \in \{100, 500, 5000\}$ . White circles represent the means.

$(X_i, Y_i)$  where  $X_i$  follows a  $d = 5$  dimensional standard normal distribution and

$$Y_i = \frac{1}{1 + \exp(-w^T X_i)} + \varepsilon_i,$$

where  $w \in \mathbb{R}^d$  follows a standard normal distribution of dimension  $d$  and  $\varepsilon_i \sim \mathcal{N}(0, 0.07)$ . The rank of the observation  $i$  is defined as the rank of the corresponding  $Y_i$ , i.e.,  $R_i = \mathbf{R}(Y_i, \{Y_j\}_{j \in [n]})$ . The calibration set is then  $\{(X_i, R_i)\}_{i \in [n]}$ . The test set is composed of the remaining  $m$  observations where only the features  $X_i$  are observed, i.e., the test set is  $\{X_{n+i}\}_{i \in [m]}$ . The sizes of the calibration and test sets are chosen among  $n, m \in \{100, 500, 2500\}$  and for each pair  $(n, m)$ , we evaluate the performance of the different methods on  $B = 1000$  data sets.

**Ranking algorithms:** We consider two well-known algorithms of different nature. The first one is RankSVM (Herbrich, Graepel, et al., 1999) which learns a score function to rank items based on their features. The ranking errors come from learning this function. The second algorithm, called the Balanced Rank Estimation algorithm (BRE, Wauthier et al., 2013), ranks items using a small number of pairwise comparisons. The ranking errors are due to this limited number of comparisons.

1. **RankSVM.** The ranking problem is first transformed into a pairwise classification task via a pairwise transformation. Formally, given  $n_{tr}$  training data initially of the form  $(X_i, R_i) \in \mathbb{R}^d \times \mathbb{N}$ , we construct the tuples  $(X_i - X_j, 2 \cdot \mathbf{1}(R_i \leq R_j) - 1)$  in  $\mathbb{R}^d \times \{-1, 1\}$  for  $i \neq j$ . The ranking problem is thus converted into a binary classification task with train sample of size  $n_{tr}(n_{tr} - 1)$ . We solve it using a linear SVM classifier with a squared hinge loss and with hyperparameters set by cross-validation. This algorithm learns a score function and ranks the points from these scores.
2. **BRE.** The algorithm randomly selects  $p$  pairs  $(i, j)$ ,  $i < j$  in  $\llbracket n + m \rrbracket^2$  and computes  $c_{i,j} = \mathbf{1}(R_i < R_j)$ . Then, for each item  $i$ , it computes the score  $\Pi(i) = \sum_{j \neq i} s_{i,j}(2c_{i,j} - 1)$  where  $s_{i,j} = 1$  if the pair  $(i, j)$  has been selected and 0 otherwise. Finally, BRE returns the estimated ranks deduced from these scores.

**Metrics.** To assess the validity of our methods, for each pair of sizes  $(n, m)$ , we collect the following empirical measurements computed on the  $B$  data sets:

1. The FCP obtained on the  $b$ -th data set:  $\text{FCP}_b$ .
2. The *relative length* of the prediction sets defined as the average size of the sets divided by  $n + m$ .

For each algorithm, we construct prediction sets using the score function  $s^{\text{RA}}$  and  $s^{\text{VA}}$  and with both the linear and quantile envelopes described in Section 4.2. The parameters  $\alpha$ ,  $1 - \beta$  and  $\delta$  equal to respectively the probability of miscoverage, the probability to control the FCP at level  $\alpha$ , and the probability of the quantile envelope, are set to 0.1, 0.75 and 0.02 respectively.

**Results.** RankSVM algorithm is learned on an independent training set of size  $n_{tr} = 300$ . Figure 2 top panel displays, for  $m = 500$  and  $n \in \{100, 500, 2500\}$ , boxplots of  $\{\text{FCP}_b\}_b$  (left) and of the relative length of the returned prediction sets (right) obtained with score functions  $s^{\text{RA}}$  or  $s^{\text{VA}}$  and with quantile or linear envelop (respectively denoted  $Qenv$  and  $Lenv$  in the legend).



## Transductive Conformal Inference for Ranking

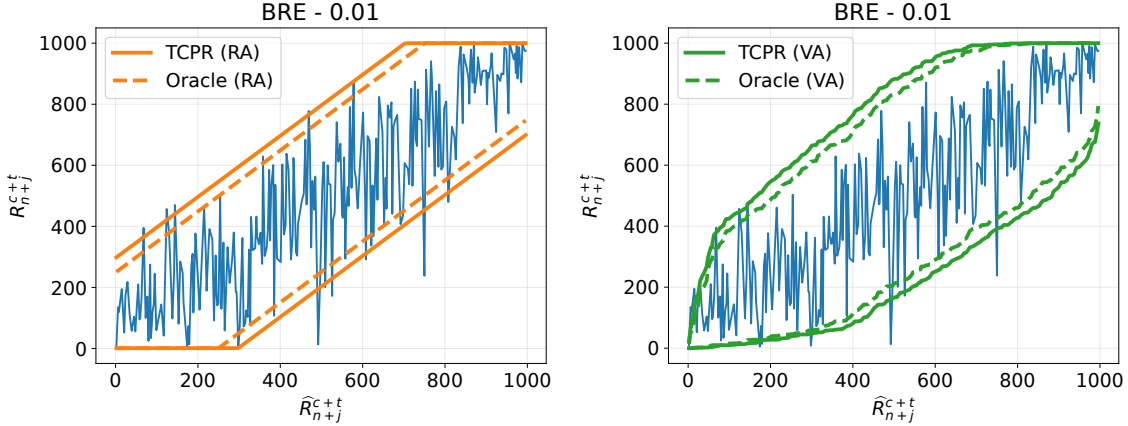


Figure 3: Synthetic data: True ranks  $R_{n+j}^{c+t}$  in function of their predicted rank  $\widehat{R}_{n+j}^{c+t}$  by BRE-0.01 and their covering sets with scores  $s^{\mathbf{RA}}$  and  $s^{\mathbf{VA}}$  for  $n = m = 500$ .

The white circle corresponds to the mean. For BRE, the number of pairwise comparisons is set to a proportion 0.01 of the total number of possible comparisons. We denote it BRE-0.01 in the sequel. The results are presented in Figure 2 bottom panel. Results for other values of  $n$  and  $m$  are presented in Appendix C.

We can first remark that the FCPs are correctly controlled in our method TCPR as they stay in general below  $\alpha = 0.1$ . They are however conservative for small values of  $n$ . This effect is due to the use of the proxy score but vanishes when the calibration size  $n$  increases. For instance, when  $n = 2500$ , our methods obtain similar performance than the oracle for all settings. In (**RA**) setting, for RankSVM and  $n = 2500$ , the length of our prediction sets are around 27% of the points against 25% for the oracle method.

The prediction sets constructed with the score  $s^{\mathbf{RA}}$  are, in average, smaller than those obtained with  $s^{\mathbf{VA}}$ . This is particularly noticeable for RankSVM. This difference is less pronounced for BRE algorithm for which both scores obtained similar performance. In fact, the size of the sets for score  $s^{\mathbf{VA}}$  varies from one point to another. In Figure 3, we observe indeed that the sets returned using  $s^{\mathbf{RA}}$  are in average smaller than those returned using  $s^{\mathbf{VA}}$ , but for small and large ranks, the sets with  $s^{\mathbf{VA}}$  are much narrower. The score  $S^{\mathbf{VA}}$  is in fact more adaptive than  $s^{\mathbf{RA}}$ , this is highlighting by another example in Section C.1 of the Appendix.

## Transductive Conformal Inference for Ranking

	FCP	Relative length	Oracle length ratio
	<b>RankSVM</b>		
<b>(RA)</b>	0.070( $\pm 0.009$ )	0.784( $\pm 0.014$ )	0.933( $\pm 0.004$ )
<b>(VA)</b>	0.061( $\pm 0.008$ )	0.672( $\pm 0.010$ )	0.930( $\pm 0.006$ )
	<b>BRE - 0.01</b>		
<b>(RA)</b>	0.044( $\pm 0.008$ )	0.421( $\pm 0.007$ )	0.875( $\pm 0.006$ )
<b>(VA)</b>	0.051( $\pm 0.007$ )	0.400( $\pm 0.007$ )	0.892( $\pm 0.005$ )
	<b>BRE - 0.1</b>		
<b>(RA)</b>	0.011( $\pm 0.003$ )	0.169( $\pm 0.002$ )	0.688( $\pm 0.011$ )
<b>(VA)</b>	0.011( $\pm 0.002$ )	0.165( $\pm 0.002$ )	0.706( $\pm 0.013$ )

Table 1: Real data: FCP, relative length of the prediction sets and ratio of the length with the oracle length for abalone data set for different algorithms and score functions.

	RankSVM	BRE-0.01	BRE-0.1
<b>NDCG</b>	0.973( $\pm 0.003$ )	0.989( $\pm 0.000$ )	0.997( $\pm 0.000$ )
<b>Rank error</b>	384.7( $\pm 4.201$ )	192.3( $\pm 2.323$ )	62.86( $\pm 1.307$ )

Table 2: Real data: NDCG and average rank error of each algorithm.

## 5.2 Real data

We now evaluate our method TCPD on the abalone data set (Nash et al., 1994). Initially, the goal of this data set is to predict the age of abalone from physical measurements. The true age of abalones is commonly determined by cutting the shell and by counting the number of rings through a microscope. Because this is a time-consuming task, predicting the age of abalone using other more accessible measurements is of huge interest. The data set is composed of the following 9 variables: sex, length, diameter, height, whole weight, shucked weight, viscera weight, shell weight and rings. The variable to predict is the age which is equal to (number of rings + 1.5). We break ties at random to obtain a full ranking.

To fit our setting, we suppose that we do not search to predict the exact age of the abalones but instead to rank the abalones them from the youngest to the oldest. We split this data set, of size 4177, into three parts: a training set of size  $n_{tr} = 1677$ , a calibration set of size  $n = 1455$  where only the true ranks inside this calibration set are kept, and a test set of size  $m = 1045$ . For the ranking task, we use (i) RankSVM with a squared hinge loss and with hyperparameters set by cross-validation and (ii) BRE where the number of pairwise comparisons is equal to a proportion 0.01 or 0.1

of the total possible number of comparisons (see previous section for details on these algorithms). We repeat this process 10 times, set  $\alpha = 0.1$ ,  $\beta = 0.2$ ,  $\delta = 0.02$ , and report the average FCP, the average relative length of the predict sets, and the ratio between the relative length obtained with our method and the one obtained with the Oracle method.

**Results.** Table 1 reports the results obtained for RankSVM and BRE with our method TCPR. As with the synthetic data, we observe that the size of the sets decreases with the performance of the algorithm. Table 1 could be compared with Table 2 which contains the Normalized Discounted Cumulative Gain (NDCG) of each algorithm and the average rank error. Conversely, as the algorithm’s performance increases, the use of proxy scores has a greater impact and the sets constructed are then further away from the oracle, as can be seen in Table 1. The ratio with the oracle increases as performance increases, although relative length decreases at the same time.

More precisely, BRE-0.1 has the higher NDCG and smaller sets than those of BRE-0.01 and even smaller than those of RankSVM. However, its FCP is also much smaller which means that the size of the prediction sets could be reduced while still maintaining a FCP smaller than 0.1. This observation is reinforced by comparing the ratios with the length of the oracle. The ratio is closer to 1 for RankSVM, which performs less well and therefore has large prediction sets. This higher ratio for RankSVM can be explained by the fact that the size of the sets depends on the error of the algorithm but also on the envelope. When the error is large, the influence of the envelope becomes negligible and the set lengths get close to the oracle. Finally, notice that for this data set, the score function  $s^{\text{VA}}$  provides smaller prediction sets. For instance, for RankSVM and for BRE-0.1, the FCP and the relative length are simultaneously smaller than those provided by the score  $s^{\text{RA}}$ .

## 6 Conclusion

We have developed a conformal prediction method that quantify the error of an algorithm ranking  $m$  new items among  $n$  previously ranked items. Our approach is based on constructing an envelope around these  $n$  ranks to quantify the algorithm’s error. Both theoretical and empirical evidence demonstrates that this envelope does not significantly impact the size of the intervals, especially when  $n$  is large relatively to  $m$ .

Since our focus has been exclusively on full ranking, it would be interesting to extend our method to a partial ranking framework. Another important line of research

would be to cover top- $k$  algorithms. Finally, the dependence of our method on the choice of the envelope we select should be investigated. We expect that a choice more adapted to the problem considered could improve the size of the prediction sets.

## Acknowledgements

The authors acknowledge Ulysse Gazin and Etienne Roquain for helpful comments. G.B. gratefully acknowledges funding from the grants ANR-21-CE23-0035 (ASCAI), ANR-19-CHIA-0021-01 (BISCOTTE) and ANR-23-CE40-0018-01 (BACKUP) of the French National Research Agency ANR. P.H. gratefully acknowledges the Emergence project MARS of Sorbonne Université.

## References

- Ailon, N. (2012). “An Active Learning Algorithm for Ranking from Pairwise Preferences with an Almost Optimal Query Complexity.” In: *Journal of Machine Learning Research* 13.1.
- Alman, J. and D. McKay (2017). “Theoretical foundations of team matchmaking”. In: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pp. 1073–1081.
- Angelopoulos, A. N., S. Bates, et al. (2023). “Conformal prediction: A gentle introduction”. In: *Foundations and Trends® in Machine Learning* 16.4, pp. 494–591.
- Angelopoulos, A. N., K. Krauth, S. Bates, Y. Wang, and M. I. Jordan (2023). “Recommendation systems with distribution-free reliability guarantees”. In: *Conformal and Probabilistic Prediction with Applications*. PMLR, pp. 175–193.
- Braverman, M. and E. Mossel (2007). “Noisy sorting without resampling”. In: *ACM-SIAM Symposium on Discrete Algorithms*.
- Braverman, M. and E. Mossel (2009). “Sorting from Noisy Information”. In: *CoRR* abs/0910.1191.
- Burges, C., T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender (2005). “Learning to rank using gradient descent”. In: *Proceedings of the 22nd international conference on Machine learning*, pp. 89–96.
- Cao, Y., J. Xu, T.-Y. Liu, H. Li, Y. Huang, and H.-W. Hon (2006). “Adapting ranking SVM to document retrieval”. In: *Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 186–193.

- Cao, Z., T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li (2007). “Learning to rank: from pairwise approach to listwise approach”. In: *Proceedings of the 24th international conference on Machine learning*, pp. 129–136.
- Feige, U., P. Raghavan, D. Peleg, and E. Upfal (1994). “Computing with noisy information”. In: *SIAM Journal on Computing* 23.5, pp. 1001–1018.
- Fontana, M., G. Zeni, and S. Vantini (2023). “Conformal prediction: a unified review of theory and new challenges”. In: *Bernoulli* 29.1, pp. 1–23.
- Freund, Y., R. Iyer, R. E. Schapire, and Y. Singer (2003). “An efficient boosting algorithm for combining preferences”. In: *Journal of machine learning research* 4.Nov, pp. 933–969.
- Gazin, U. (2024). “Asymptotics for conformal inference”. In: *arXiv preprint arXiv:2409.12019*.
- Gazin, U., G. Blanchard, and E. Roquain (2024). “Transductive conformal inference with adaptive scores”. In: *International Conference on Artificial Intelligence and Statistics*. PMLR, pp. 1504–1512.
- Heckel, R., N. B. Shah, K. Ramchandran, and M. J. Wainwright (2019). “Active ranking from pairwise comparisons and when parametric assumptions do not help”. In: *The Annals of Statistics* 47.6, pp. 3099–3126.
- Herbrich, R., T. Graepel, and K. Obermayer (1999). “Support vector learning for ordinal regression”. In: *1999 Ninth International Conference on Artificial Neural Networks ICANN 99. (Conf. Publ. No. 470)*. Vol. 1, pp. 97–102.
- Herbrich, R., T. Minka, and T. Graepel (2006). “TrueSkill™: a Bayesian skill rating system”. In: *Advances in neural information processing systems* 19.
- Huang, K., Y. Jin, E. Candes, and J. Leskovec (2024). “Uncertainty quantification over graph with conformalized graph neural networks”. In: *Advances in Neural Information Processing Systems* 36.
- Lei, J., M. G’Sell, A. Rinaldo, R. J. Tibshirani, and L. Wasserman (2018). “Distribution-free predictive inference for regression”. In: *Journal of the American Statistical Association* 113.523, pp. 1094–1111.
- Li, H. (2011). “A short introduction to learning to rank”. In: *IEICE TRANSACTIONS on Information and Systems* 94.10, pp. 1854–1862.
- Liu, N. N. and Q. Yang (2008). “Eigenrank: a ranking-oriented approach to collaborative filtering”. In: *Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 83–90.
- Liu, T.-Y. et al. (2009). “Learning to rank for information retrieval”. In: *Foundations and Trends® in Information Retrieval* 3.3, pp. 225–331.
- Liu, Y., J.-W. Bi, and Z.-P. Fan (2017). “Ranking products through online reviews: A method based on sentiment analysis technique and intuitionistic fuzzy set theory”. In: *Information Fusion* 36, pp. 149–161.

- Manokhin, V. (2024). *Awesome Conformal Prediction*. <https://github.com/valeman/awesome-conformal-prediction>. Github repository.
- Marques F., P. C. (2024). “Universal distribution of the empirical coverage in split conformal prediction”. In: *arXiv preprint*. arXiv: [2303.02770](https://arxiv.org/abs/2303.02770).
- Minka, T., R. Cleven, and Y. Zaykov (2018). “Trueskill 2: An improved bayesian skill rating system”. In: *Technical Report*.
- Naaman, M. (2021). “On the tight constant in the multivariate Dvoretzky–Kiefer–Wolfowitz inequality”. In: *Statistics & Probability Letters* 173, p. 109088.
- Nash, W., T. Sellers, S. Talbot, A. Cawthorn, and W. Ford (1994). *Abalone*. UCI Machine Learning Repository. DOI: <https://doi.org/10.24432/C55C7W>.
- Negahban, S., S. Oh, and D. Shah (2012). “Iterative ranking from pair-wise comparisons”. In: *Advances in neural information processing systems* 25.
- Papadopoulos, H., K. Proedrou, V. Vovk, and A. Gammerman (2002). “Inductive confidence machines for regression”. In: *Machine Learning: ECML 2002: 13th European Conference on Machine Learning Helsinki, Finland, August 19–23, 2002 Proceedings 13*. Springer, pp. 345–356.
- Pedregosa, F., E. Cauvet, G. Varoquaux, C. Pallier, B. Thirion, and A. Gramfort (2012). “Learning to rank from medical imaging data”. In: *Machine Learning in Medical Imaging: Third International Workshop, MLMI 2012, Held in Conjunction with MICCAI 2012, Nice, France, October 1, 2012, Revised Selected Papers 3*. Springer, pp. 234–241.
- Romano, Y., E. Patterson, and E. Candes (2019). “Conformalized quantile regression”. In: *Advances in neural information processing systems* 32.
- Vovk, V. (2013). “Transductive conformal predictors”. In: *Artificial Intelligence Applications and Innovations: 9th IFIP WG 12.5 International Conference, AIAI 2013, Paphos, Cyprus, September 30–October 2, 2013, Proceedings 9*. Springer, pp. 348–360.
- Vovk, V., A. Gammerman, and G. Shafer (2005). *Algorithmic learning in a random world*. Springer Science & Business Media.
- Wauthier, F., M. Jordan, and N. Jojic (2013). “Efficient ranking from pairwise comparisons”. In: *International Conference on Machine Learning*. PMLR, pp. 109–117.
- Xu, Y., W. Guo, and Z. Wei (2024). “Conformal Ranked Retrieval”. In: *arXiv preprint arXiv:2404.17769*.

## A Proofs

### A.1 Proof of Theorem 3

Let us denote  $E_\delta = \{\forall i \in \llbracket n \rrbracket : R_i^{c+t} \in [R_i^-, R_i^+]\}$ . On the event  $E_\delta$ , for all  $i \in \llbracket n \rrbracket$ , we have:

$$S_i^{\text{true}} := s(X_i, R_i^{c+t}, (X_k)_{k \in \llbracket n+m \rrbracket}) \leq \max_{r \in [R_i^-, R_i^+]} s(X_i, r, (X_k)_{k \in \llbracket n+m \rrbracket}) := S_i.$$

Hence, the quantiles of the “true” scores  $S_i^{\text{true}}$  are upper bounded by the quantiles of the proxy scores  $\{S_i\}_{i \in \llbracket n \rrbracket}$ , i.e., for all  $k \in \llbracket n \rrbracket$ :

$$S_{(k)}^{\text{true}} \leq S_{(k)}. \quad (12)$$

Then, for  $i \in \llbracket m \rrbracket$ :

$$\begin{aligned} \mathbb{P}\left(R_{n+i}^{c+t} \notin \widehat{\mathcal{C}}(X_{n+i})\right) &\leq \mathbb{P}\left(R_{n+i}^{c+t} \notin \widehat{\mathcal{C}}(X_{n+i}) \cap E_\delta\right) + \mathbb{P}(E_\delta^c) \\ &\leq \mathbb{P}\left(s(X_{n+i}, R_{n+i}^{c+t}, (X_j)_{j \in \llbracket n+m \rrbracket}) > S_{(k)} \cap E_\delta\right) + \mathbb{P}(E_\delta^c) \\ &\leq \mathbb{P}\left(s(X_{n+i}, R_{n+i}^{c+t}, (X_j)_{j \in \llbracket n+m \rrbracket}) > S_{(k)}^{\text{true}}\right) + \delta \leq \alpha - \delta + \delta = \alpha. \end{aligned}$$

To pass from the second to the third line, we have used (15) which is satisfied on the event  $E_\delta$ . The last inequality is obtained using that if the score function  $s(x, r, (x_j)_j)$  is invariant by permutation of its last argument and the vectors  $(X_i, R_i^{c+t})_{i \in \llbracket n+m \rrbracket}$  are exchangeable, then the true scores  $\{S_i^{\text{true}}\}_{i \in \llbracket n \rrbracket}$  are exchangeable. Therefore, we can use Theorem 1 with  $k = \lceil (n+1)(1-\alpha+\delta) \rceil$ .

### A.2 Restatement of Gazin et al. (2024, Theorem 2.4)

We restate here Theorem 2.4 of Gazin et al. (2024) which controls the empirical cumulative distribution function of conformal  $p$ -values.

**Theorem 11.** *Let  $n, m \geq 2$ ,  $V_1, \dots, V_{n+m}$  be exchangeable real random variables having no ties almost surely. Let us, for  $1 \leq i \leq n$ , define the conformal  $p$ -values:*

$$p_i = \frac{1}{m+1} \left( 1 + \sum_{j=1}^m \mathbf{1}\{V_i \geq V_{n+j}\} \right) \quad (13)$$

## Transductive Conformal Inference for Ranking

and  $\widehat{F}_n(t) := n^{-1} \sum_{i=1}^n \mathbf{1}\{p_i \leq t\}$  be their empirical cumulative distribution function. Then, for  $\beta \in (0, 1)$ :

$$\mathbb{P} \left[ \sup_{t \in [0,1]} \left| \widehat{F}_n(t) - I_m(t) \right| > \sqrt{\frac{\log(1/\beta) + \log(1 + \sqrt{2\pi\tau_{n,m}})}{2\tau_{n,m}}} \right] \leq 2\beta, \quad (14)$$

where  $\tau_{n,m} = \frac{nm}{n+m}$  and  $I_m(t) = \lfloor (m+1)t \rfloor / (m+1)$ .

*Proof.* The concentration inequality is obtained from Gazin et al. (2024, Theorem 2.4). The lower deviation was enounced in its proof (Equation (38)). We apply the theorem with  $r = 1$  and lower bound  $\tau_{n,m}/\sqrt{n+m}$  by  $\sqrt{\tau_{n,m}}/2$ .  $\square$

### A.3 Proof of Proposition 1

On the event  $\{\forall i \in \llbracket n \rrbracket : R_i^{c+t} \in [R_i^-, R_i^+]\}$ , we have:

$$S_i^{\text{true}} := s(X_i, R_i^{c+t}, (X_k)_{k \in \llbracket n+m \rrbracket}) \leq \max_{r \in [R_i^-, R_i^+]} s(X_i, r, (X_k)_{k \in \llbracket n+m \rrbracket}) := S_i.$$

Hence, as earlier the quantiles of the ‘‘true’’ scores  $S_i^{\text{true}}$  are upper bounded by the quantiles of the proxy scores  $\{S_i\}_{i \in \llbracket n \rrbracket}$ , i.e., for all  $k \in \llbracket n \rrbracket$ :

$$S_{(k)}^{\text{true}} \leq S_{(k)}. \quad (15)$$

Then, on this event:

$$\sum_{i=1}^m \mathbf{1}\{R_{n+i}^{c+t} \notin \widehat{C}_k(X_{n+i})\} = \sum_{i=1}^m \mathbf{1}\{S_{n+i}^{\text{true}} > S_{(k)}\} \leq \sum_{i=1}^m \mathbf{1}\{S_{n+i}^{\text{true}} > S_{(k)}^{\text{true}}\}. \quad (16)$$

If the true scores have no ties, we can directly apply Theorem 11 to obtain a control of the FCP. The following technical details allow to control it in the case where the score function is discrete, for instance when  $s = s^{\text{RA}}$ . Let us introduce

$$\Gamma := \min_{\substack{1 \leq i, j \leq n+m \\ S_i^{\text{true}} \neq S_j^{\text{true}}}} |S_i^{\text{true}} - S_j^{\text{true}}|$$

and an independent family of random variables  $\gamma_i$  for  $i \in \llbracket n+m \rrbracket$  (also independent of  $(S_i^{\text{true}})_{i \in \llbracket n+m \rrbracket}$ ) following a uniform distribution on  $(-\Gamma/2, \Gamma/2)$ . Define  $\widetilde{S}_i := S_i^{\text{true}} + \gamma_i$  for  $i \in \llbracket n+m \rrbracket$ . Observe that  $\Gamma$  depends on the true scores, but is invariant by



## Transductive Conformal Inference for Ranking

permutation of these scores. Hence, the ‘‘perturbed’’ scores  $\tilde{S}_i$  are exchangeable and have (almost surely) no ties<sup>1</sup>. Moreover for  $i \in \llbracket m \rrbracket$ :

$$\begin{aligned} \{S_{n+i}^{\text{true}} > S_{(k)}^{\text{true}}\} &= \left\{ \sum_{j=1}^n \mathbf{1}\{S_{n+i}^{\text{true}} > S_j^{\text{true}}\} \geq k \right\} \\ &\subseteq \left\{ \sum_{j=1}^n \mathbf{1}\{S_{n+i}^{\text{true}} + \gamma_{n+i} > S_j^{\text{true}} + \gamma_j\} \geq k \right\} \\ &= \left\{ \sum_{j=1}^n \mathbf{1}\{\tilde{S}_{n+i} > \tilde{S}_j\} \geq k \right\} = \{\tilde{S}_{n+i} > \tilde{S}_{(k)}\}. \end{aligned}$$

The main argument of these inequalities is that by construction of the random variables  $\gamma_i$ , for all  $i, j$ , the event  $\{S_i^{\text{true}} > S_j^{\text{true}}\}$  is included in  $\{\tilde{S}_i > \tilde{S}_j\}$ .

Then we can upper bound the FCP using (16):

$$\frac{1}{m} \sum_{i=1}^m \mathbf{1}\{R_{n+i}^{c+t} \notin \hat{\mathcal{C}}_k(X_{n+i})\} \leq \frac{1}{m} \sum_{i=1}^m \mathbf{1}\{\tilde{S}_{n+i} > \tilde{S}_{(k)}\}. \quad (17)$$

We want to apply Theorem 11 to the exchangeable random variables  $\tilde{S}$ , which have no ties. Let us express the FCP in term of conformal  $p$ -values. We introduce for  $i \in \llbracket m \rrbracket$

$$\tilde{p}_i := \frac{1}{n+1} \left( 1 + \sum_{j=1}^n \mathbf{1}\{\tilde{S}_{n+i} \geq \tilde{S}_j\} \right)$$

and for  $t \in (0, 1)$

$$\tilde{F}_m(t) := \frac{1}{m} \sum_{i=1}^m \mathbf{1}\{\tilde{p}_i \leq t\}.$$

Then for  $i \in \llbracket m \rrbracket$ , the following events are equal:

$$\begin{aligned} \{\tilde{S}_{n+i} > \tilde{S}_{(k)}\} &= \left\{ \sum_{j=1}^n \mathbf{1}\{\tilde{S}_{n+i} > \tilde{S}_j\} \geq k \right\} = \left\{ \sum_{j=1}^n \mathbf{1}\{\tilde{S}_{n+i} \geq \tilde{S}_j\} \geq k \right\} \\ &= \left\{ \sum_{j=1}^n \mathbf{1}\{\tilde{S}_{n+i} \geq \tilde{S}_j\} > k - 1 \right\} = \{\tilde{p}_i > k/(n+1)\}, \end{aligned}$$

---

<sup>1</sup> $\Gamma$  is not defined if all scores are equal. However, in this case the FCP is null and then upper bounded by  $\alpha$ .

where we have used that the scores  $\tilde{S}$  have no ties. Injecting these equalities into (17), we get:

$$\frac{1}{m} \sum_{i=1}^m \mathbf{1}\{R_{n+i}^{c+t} \notin \hat{\mathcal{C}}_k(X_{n+i})\} \leq 1 - \tilde{F}_m(k/(n+1)).$$

Then, with probability  $1 - (\beta + \delta)$ , using (14) (inverting  $n$  and  $m$ ), we get:

$$\begin{aligned} \frac{1}{m} \sum_{i=1}^m \mathbf{1}\{R_{n+i}^{c+t} \notin \hat{\mathcal{C}}_k(X_{n+i})\} &\leq 1 - I_n(k/(n+1)) + \lambda_{n,m} \\ &= 1 - \frac{k}{n+1} + \lambda_{n,m} \leq \alpha + \lambda_{n,m}, \end{aligned}$$

where  $\lambda_{n,m} = \sqrt{\frac{\log(C\sqrt{\tau_{n,m}}/\beta)}{\tau_{n,m}}}$ . □

#### A.4 Proof of Proposition 2

Let us apply Theorem 11 to  $V_i = Y_i$  for  $1 \leq i \leq n+m$ . Then the  $p$ -value  $p_i$ -s are affine transformations of the ranks of the calibration point in the test sample:

$$p_i = \frac{1}{m+1}(1 + R_i^t),$$

where we recall that  $R_i^t := \mathbf{R}(Y_i, \{Y_{n+j}\}_{j \in \llbracket m \rrbracket})$ . Let us then rewrite the empirical cumulative distribution function  $\hat{F}_n$  in function of  $R_i^t$ , for  $t \in (0, 1)$ :

$$\hat{F}_n(t) = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{R_i^t \leq (m+1)t - 1\}.$$

Let us now remark that for some  $u \in \mathbb{R}_+$ , it is equivalent for a calibration sample point  $j$  to have its rank in the test set smaller than  $u$  than to have more points than its rank in the calibration with a rank smaller than  $u$ :

$$\forall j \in \llbracket n \rrbracket : \quad R_j^t \leq u \iff \sum_{i=1}^n \mathbf{1}\{R_i^t \leq u\} \geq R_j^c \iff n\hat{F}_n((u+1)/(m+1)) \geq R_j^c.$$

We point out that the first equivalence is satisfied as the order of the items in the calibration set is conserved when considering their insertion ranks in the test set: for

## Transductive Conformal Inference for Ranking

all  $1 \leq i, j \leq n$ ,  $R_i^c \leq R_j^c \Rightarrow R_i^t \leq R_j^t$ . Thus, for  $j \in \llbracket n \rrbracket$ , if  $R_j^t \leq u$ , the  $R_j^c$  calibration items  $i \in \llbracket n \rrbracket$  satisfying  $R_i^c \leq R_j^c$  also satisfy  $R_i^t \leq u$ , showing the direct implication ( $\Rightarrow$ ) in the last display. Conversely, if  $k$  is the number of calibration items  $i \in \llbracket n \rrbracket$  satisfying  $R_i^t \leq u$ , they are necessarily exactly the items corresponding to the  $k$  smallest calibration ranks. Otherwise, there would be a calibration item  $j \in \llbracket n \rrbracket$  with  $R_j^c = k' > k$  and  $R_j^t \leq u$ , implying (by the previous token) that  $\sum_{i=1}^n \mathbf{1}\{R_i^t \leq u\} \geq k' > k$ , a contradiction. Thus any calibration item with  $R_j^c \leq k$  satisfies  $R_j^t \leq u$ .

Thus, with probability  $1 - 2\delta$ , thanks to (14), for all  $t \in (0, 1)$ :

$$\left| \widehat{F}_n(t) - I_m(t) \right| \leq \lambda, \quad \text{where} \quad \lambda = \sqrt{\frac{\log(1/\delta) + \log(1 + \sqrt{2\pi\tau_{n,m}})}{2\tau_{n,m}}}.$$

Assume this event is satisfied. Then in particular, for  $u = (m + 1)(R_j^c/n + \lambda)$ :

$$\begin{aligned} n\widehat{F}_n((u + 1)/(m + 1)) &= n\widehat{F}_n(R_j^c/n + \lambda + 1/(m + 1)) \\ &\geq nI_m(R_j^c/n + \lambda + 1/(m + 1)) - n\lambda \\ &= n\lfloor (m + 1)(R_j^c/n + \lambda) + 1 \rfloor / (m + 1) - n\lambda \\ &\geq n(R_j^c/n + \lambda) - n\lambda \geq R_j^c. \end{aligned}$$

Thus,  $R_j^t \leq (m + 1)(R_j^c/n + \lambda)$  for  $1 \leq j \leq n$ . Let us now lower bound these ranks. Similarly, we have the equivalence for  $u \in \mathbb{R}$  and  $1 \leq j \leq n$ :

$$R_j^t > u \iff \sum_{i=1}^n \mathbf{1}\{R_i^t \leq u\} \leq R_j^c \iff n\widehat{F}_n((u + 1)/(m + 1)) \leq R_j^c.$$

With  $u = (m + 1)(R_j^c/n - \lambda)$ , we have  $n\widehat{F}_n((u + 1)/(m + 1)) \leq R_j^c$  and then  $R_j^t \geq (m + 1)(R_j^c/n - \lambda)$  for  $1 \leq j \leq n$  which concludes the proof.  $\square$

### A.5 Proof of Proposition 10

The proof is based on the concentration bound of Naaman (2021) recalled below which is a recent multi-dimensional version of the Dvoretzky–Kiefer–Wolfowitz (DKW) inequality.

**Theorem 12** (Naaman, 2021). *Let  $X^{(1)}, \dots, X^{(K)}$  i.i.d. random vectors in  $\mathbb{R}^n$ . Then for  $t \geq 0$ ,*

$$\mathbb{P} \left[ \sup_{\theta \in \mathbb{R}^n} \left| F(\theta) - \widehat{F}_K(\theta) \right| \geq t \right] \leq n(K + 1)e^{-2Kt^2}, \quad (18)$$

## Transductive Conformal Inference for Ranking

where  $F(\theta) = \mathbb{P}\left[X_i^{(1)} \leq \theta_i, \forall i\right]$  and  $\widehat{F}_K$  is the empirical cdf defined by:

$$\widehat{F}_K(\theta) = \frac{1}{K} \sum_{k=1}^K \mathbf{1}\left\{X_i^{(k)} \leq \theta_i, \forall i\right\}.$$

**Proof of Proposition 10.** For  $\theta \in \mathbb{R}^n$ , let us denote

$$F(\theta) = \mathbb{P}\left[\forall i \in \llbracket n \rrbracket : R_{(i)}^{c+t} \leq \theta_i\right]$$

and  $\widehat{F}_k$ , the empirical cdf:

$$\widehat{F}_K(\theta) = \frac{1}{K} \sum_{k=1}^K \mathbf{1}\left\{R_i^{(k)} \leq \theta_i, \forall i \in \llbracket n \rrbracket\right\}.$$

Let us now remark that:

$$\begin{aligned} \mathbb{P}\left[\forall i \in \llbracket n \rrbracket : R_i^{c+t} \in \left[R_{R_i^-}^-, R_{R_i^+}^+\right]\right] &= \mathbb{P}\left[\forall i \in \llbracket n \rrbracket : R_{(i)}^{c+t} \in \left[R_i^-, R_i^+\right]\right] \\ &= \mathbb{E}\left[F(R^+) - F(R^-)\right]. \end{aligned}$$

The ordered ranking vector  $R_{(i)}^{c+t}$  follows indeed the same distribution than the drawn vectors  $R_i^{(k)}$ . Then:

$$\begin{aligned} F(R^+) - F(R^-) &= \left(F(R^+) - \widehat{F}_K(R^+)\right) - \left(F(R^-) - \widehat{F}_K(R^-)\right) + \widehat{F}_K(R^+) - \widehat{F}_K(R^-) \\ &\geq -2 \sup_{\theta \in \mathbb{R}^n} \left|F(\theta) - \widehat{F}_K(\theta)\right| + (1 - \delta), \end{aligned}$$

as by assumption on the envelope,  $\widehat{F}_K(R^+) - \widehat{F}_K(R^-) \geq 1 - \delta$  almost surely. Then after taking the expectation, we get:

$$\mathbb{P}\left[\forall i \in \llbracket n \rrbracket : R_i^{c+t} \in \left[R_{R_i^-}^-, R_{R_i^+}^+\right]\right] \geq 1 - \delta - 2\mathbb{E}\left[\sup_{\theta} \left|F(\theta) - \widehat{F}_K(\theta)\right|\right].$$

It remains to upper bound that expectation. Let us denote  $Z = \sup_{\theta} \left|F(\theta) - \widehat{F}_K(\theta)\right|$ , let  $t \geq 0$ ,

$$\mathbb{E}[Z] = \mathbb{E}[Z\mathbf{1}_{Z \leq t} + Z\mathbf{1}_{Z > t}] \leq t + \mathbb{P}[Z > t] \leq t + n(K+1)e^{-2Kt^2}.$$

We have used that  $Z$  is bounded by 1 and Theorem 12. We conclude by choosing  $t = \sqrt{\frac{\log nK}{K}}$ .  $\square$

## B Detailed procedure to control the FCP

In this section, we give details on the procedure to control the FCP at level  $\bar{\alpha} \in (0, 1)$ . First, let us remark that  $\text{FCP} := m^{-1} \sum_{i=1}^m \mathbf{1}\{p_i \leq t\}$ , where the  $p_i$  are given in Eq. (13). In other words, the FCP is the proportion of conformal p-values that are lower than  $t$ . Furthermore, we have that (Gazin et al., 2024):

$$\begin{aligned} \{\text{FCP} \leq \bar{\alpha}\} &= \left\{ m^{-1} \sum_{i=1}^m \mathbf{1}\{p_i \leq t\} \leq \bar{\alpha} \right\} \\ &= \left\{ \sum_{i=1}^m \mathbf{1}\{p_i \leq t\} \leq \lfloor m \cdot \bar{\alpha} \rfloor \right\} = \{p_{(\lfloor m \cdot \bar{\alpha} \rfloor + 1)} > t\}, \end{aligned}$$

where  $p_{(a)}$  denotes the  $a$ -th smallest value among  $(p_1, \dots, p_m)$ . We see that controlling the FCP at level  $\bar{\alpha}$  is equivalent to control that the ordered  $p$ -value  $p_{(\lfloor m \cdot \bar{\alpha} \rfloor + 1)}$  is strictly larger than  $t$ . Hence, for  $\bar{\beta} > 0$ , if we denote by  $t_{\bar{\beta}}$  the quantile of level  $\bar{\beta}$  of the distribution of  $p_{(\lfloor m \cdot \bar{\alpha} \rfloor + 1)}$ , then  $\mathbb{P}(\text{FCP} \leq \bar{\alpha}) = \mathbb{P}(p_{(\lfloor m \cdot \bar{\alpha} \rfloor + 1)} > t_{\bar{\beta}}) \geq 1 - \bar{\beta}$ .

It remains to find  $t_{\bar{\beta}}$ . As proven by Gazin et al. (2024, Proposition 2.1), if the scores  $V_1, \dots, V_{n+m}$  are exchangeable random variables having no ties almost surely, the vector of conformal p-values  $(p_1, \dots, p_m)$  follows a known universal distribution. Hence, it is possible to simulate samples following the same distribution as  $p_{(\lfloor m \cdot \bar{\alpha} \rfloor + 1)}$  by first simulating the vector  $(p_1, \dots, p_m)$  and then by tacking the  $(\lfloor m \cdot \bar{\alpha} \rfloor + 1)$  smallest value. Then, we compute the empirical quantile of order  $\bar{\beta}$  to find  $t_{\bar{\beta}}$ . The full procedure is given in Algorithm 3.

## C Additional experiments on synthetic data

### C.1 Adaptability of score $s^{\text{VA}}$

To further highlight the adaptability of score  $s^{\text{VA}}$  relatively to  $s^{\text{RA}}$ , we consider another data set of size  $n = m = 500$ . Each pair  $(X_i, Y_i)$  is defined by  $Y_i = X_i + \varepsilon_i$  where  $X_i \sim \text{Beta}(.04, .04)$  and  $\varepsilon_i \sim \mathcal{N}(0, .07)$ . Here, because the random variables  $(X_i)_i$  follow a Beta distribution with parameters lower than 1, the probability of observing a value near 0 or 1 is much larger than the probability of observing a value near 0.5. As there is more observations close to these values, RankSVM ranks more difficultly the observations near 0 or 1 than the others. However, as the prediction sets constructed with  $s^{\text{RA}}$  have the same size, they do not reflect this difficulty (see an example in

---

**Algorithm 3** Control of the FCP

---

- 1: **Input:**  $\bar{\alpha}, \bar{\beta}, n, m$  and  $K$
  - 2:  $\tilde{P} \leftarrow$  vector of size  $K$
  - 3: **for**  $k = 1, \dots, K$  **do**
  - 4:     Draw  $n + m$  uniform random variable  $U_i$  on  $[0, 1]$
  - 5:     **for**  $i = 1, \dots, m$  **do**
  - 6:          $\tilde{p}_i \leftarrow \frac{1}{m+1} \left( 1 + \sum_{j=1}^m \mathbf{1}\{U_{j+n} \geq U_i\} \right)$
  - 7:     **end for**
  - 8:      $\tilde{P}_k \leftarrow \hat{Q}_{\lfloor m \cdot \bar{\alpha} \rfloor + 1}(\tilde{p}_1, \dots, \tilde{p}_m)$
  - 9: **end for**
  - 10:  $\hat{t}_{\bar{\beta}} \leftarrow \hat{Q}_{\bar{\beta}}(\tilde{P})$
  - 11: **Output:**  $\hat{t}_{\bar{\beta}}$
- 

Figure 4 left panel). On the contrary, the sets returned with  $s^{\text{VA}}$  are narrower when the ranking is easier (close to  $500 = 0.5 \cdot (n + m)$ ) and wider for small and large ranks (Figure 4 right panel).

## C.2 Other results

In this section, we display results of RankSVM and BRE-0.01 for the 2 different CP methods and for  $m, n \in \{100, 500, 2500\}$ . Figures 5 and 6 display boxplots of the false coverage proportion  $\{\text{FCP}_b\}_{b \in \llbracket B \rrbracket}$  (left) for  $B = 1000$  and of the relative length of the returned prediction sets (right) obtained with score functions  $s^{\text{RA}}$  and  $s^{\text{VA}}$ .

As already observed in the main paper, our method, TCPR, always control the false coverage proportion at level  $\alpha = 0.1$ . Furthermore, at any given size  $m$ , as the size  $n$  of the calibration set increases, our methods return better sets with FCP closer to 0.1 and increasingly smaller sizes. In general, we need more calibration points than the number of test points to be able to construct prediction sets that are not too large, which is quite intuitive. For  $n$  large relatively to  $m$ , we reach the performance of the oracle. For RankSVM, the score function  $s^{\text{RA}}$  gives in general better prediction sets than those constructed with  $s^{\text{VA}}$ . However, this difference is less notable for BRE. Overall, the difference of performance between the score functions depends on the distribution of the data and the algorithm. For instance, for real data experiments,  $s^{\text{VA}}$  performs better (see Table 1).

## Transductive Conformal Inference for Ranking




Figure 4: Synthetic data: True ranks  $R_{n+j}^{c+t}$  in function of their predicted rank  $\hat{R}_{n+j}^{c+t}$  by RankSVM and their covering sets with scores  $s^{\text{RA}}$  and  $s^{\text{VA}}$  for  $n = m = 500$ .

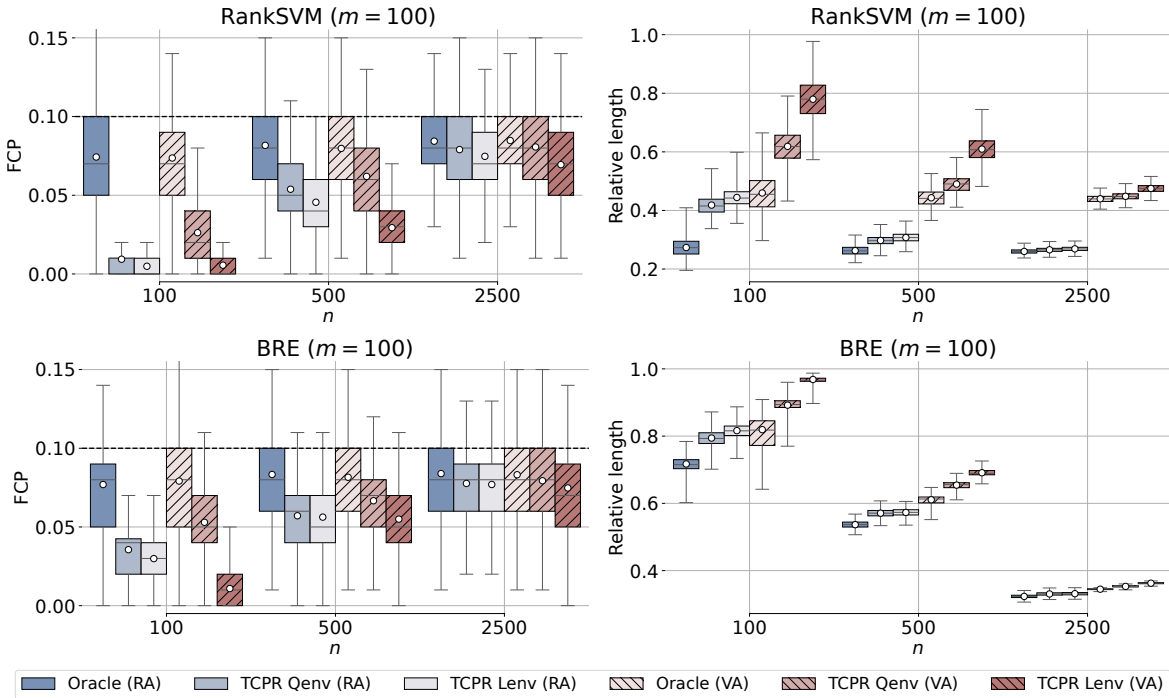


Figure 5: Synthetic data: FCP and relative lengths obtained for RankSVM and BRE-0.01 on the **(RA)** and **(VA)** settings when  $m = 100$  and  $n \in \{100, 500, 5000\}$ . White circles represent the means.

# Transductive Conformal Inference for Ranking




Figure 6: Synthetic data: FCP and relative lengths obtained for RankSVM and BRE-0.01 on the **(RA)** and **(VA)** settings when  $m = 2500$  and  $n \in \{100, 500, 5000\}$ . White circles represent the means.