
Direct Learning to Rank and Rerank: Supplementary File

AISTATS 2018

Cynthia Rudin and Yining Wang

Abstract Learning-to-rank techniques have proven to be extremely useful for prioritization problems, where we rank items in order of their estimated probabilities, and dedicate our limited resources to the top-ranked items. This work exposes a serious problem with the state of learning-to-rank algorithms, which is that they are based on convex proxies that lead to poor approximations. We then discuss the possibility of “exact” reranking algorithms based on mathematical programming. We prove that a relaxed version of the “exact” problem has the same optimal solution, and provide an empirical analysis.

Keywords learning to rank, reranking, supervised ranking, mixed-integer programming, rank statistics, discounted cumulative gain, preference learning

1 Introduction

We are often faced with prioritization problems – how can we rank aircraft in order of vulnerability to failure? How can we rank patients in order of priority for treatment? When we have limited resources and need to make decisions on how to allocate them, these ranking problems become important. The quality of a ranked list is often evaluated in terms of *rank statistics*. The area under the receiver operator characteristic curve (AUC, Metz, 1978; Bradley, 1997), which counts pairwise comparisons, is a rank statistic, but it does not focus on the top of a ranked list, and is not a good evaluation measure if we care about prioritization problems. For prioritization problems, we would use rank statistics that focus on the top of the ranked list, such as a weighted area under the curve that focuses on the left part of the curve. Then, since we evaluate our models using these rank statistics, we should aim to optimize them out-of-sample by optimizing them in-sample. The learning-to-rank field (also called supervised ranking) is built from this fundamental idea. Learning-to-rank is a natural fit for many prioritization problems. If we are able to improve the quality of a prioritization policy by even a small amount, it can have an important practical impact. Learning-to-rank can be used to prioritize mechanical equipment for repair (e.g., airplanes, as considered by Oza et al, 2009), it could be useful for prioritizing maintenance on the power grid (Rudin et al, 2012, 2010), it could be used for ranking medical workers in order of likelihood that they accessed medical records inappropriately (as considered by Menon et al, 2013), prioritizing safety inspections or lead paint inspections in dwellings (Potash et al, 2015), ranking companies in order of likeliness of committing tax violations (see Kong and Saar-Tsechansky, 2013), ranking water pipes in order of vulnerability (as considered by Li et al, 2013), other areas of information retrieval (Xu, 2007; Cao et al, 2007; Matveeva et al, 2006; Lafferty and Zhai, 2001; Li et al, 2007) and in almost any domain where one measures the quality of results by rank statistics. Learning-to-rank algorithms have been used also in sentiment analysis (Kessler and Nicolov, 2009), natural language processing (Ji et al, 2006; Collins and Koo, 2005), image retrieval (Jain and Varma, 2011; Kang et al, 2011), and reverse-engineering product quality rating systems (Chang et al, 2012).

This work exposes a serious problem with the state of learning-to-rank algorithms, which is that they are based on convex proxies for rank statistics, and when these convex proxies are used, computation is faster but the quality of the solution can be poor.

We then discuss the possibility of more direct optimization of rank statistics for predictive learning-to-rank problems. In particular, we consider a strategy of ranking with a simple ranker (logistic regression for instance) which is computationally efficient, and then reranking only the candidates near the top of the ranked list with an “exact” method. The exact method does not have the shortcoming that we discussed earlier for convex proxies.

For most ranking applications, we care only about the top of the ranked list; thus, as long as we rerank enough items with the exact method, the re-ranked list is (for practical purposes) just as useful as a full ranked list would be (if we could compute it with the exact method, which would be computationally prohibitive).

The best known theoretical guarantee on ranking methods is obtained by directly optimizing the rank statistic of interest (as shown by theoretical bounds of Clemençon and Vayatis, 2008; Rudin and Schapire, 2009, for instance) hence our choice of methodology – mixed-integer programming (MIP) – for reranking in this work. Our general formulation can optimize any member of a large class rank statistics using a single mixed-integer *linear* program. Specifically, we can handle (a generalization of) the large class of *conditional linear rank statistics*, which includes the Wilcoxon-Mann Whitney U statistic, or equivalently the Area Under the ROC Curve, the Winner-Take-All statistic, the Discounted Cumulative Gain used in information retrieval (Järvelin and Kekäläinen, 2000), and the Mean Reciprocal Rank.

Exact learning-to-rank computations need to be performed carefully; we should not refrain from solving hard problems, but certain problems are harder than others. We provide two MIP formulations aimed at the same ranking problems. The

first one works no matter what the properties of the data are. The second formulation is much faster, and is theoretically shown to produce the *same* quality of result as the first formulation when there are no duplicated observations. Note that if the observations are chosen from a continuous distribution then duplicated observations do not occur, with probability one.

One challenge in the exact learning-to-rank formulation is the way of handling ties in score. As it turns out, the original definition of conditional linear rank statistics can be used for the purpose of evaluation but not optimization. We show that a small change to the definition can be used for optimization.

This paper differs from our earlier technical report and non-archival conference paper (Chang et al, 2011, 2010), which were focused on solving full problems to optimality, and did not consider reranking or regularization; our exposition for the formulations closely follows this past work. The technique was used by Chang et al (2012) for the purpose of reverse engineering product rankings from rating companies that do not reveal their secret rating formula.

Section 2 of this paper introduces ranking and reranking, introduces the class of conditional linear rank statistics that we work with, and provides background on some current approximate algorithms for learning-to-rank. It also provides an example to show how ranked statistics can be “washed out” when they are approximated by convex substitutes. Section 2 also discusses a major difference between approximation methods and exact methods for optimizing rank statistics, which is how to handle ties in rank. As it turns out, we cannot optimize conditional linear rank statistics without changing their definition: a tie in score needs to be counted as a mistake. Section 3 provides the two MIP formulations for ranking, and Section 4 contains a proof that the second formulation is sufficient to solve the ranking problem provided that no observations are duplicates of each other. Then follows an empirical discussion in Section 5, designed to highlight the tradeoffs in the quality of the solution outlined above. Appendix A.1 contains a MIP formulation for regularized AUC maximization, and Appendix A.2 contains a MIP formulation for a general (non bipartite) ranking problem.

The recent work most related to ours are possibly those of Ataman et al (2006) who proposed a ranking algorithm to maximize the AUC using linear programming, and Brooks (2010), who uses a ramp loss and hard margin loss rather than a conventional hinge loss, making their method robust to outliers, within a mixed-integer programming framework. The work of Tan et al (2013) uses a non-mathematical-programming coordinate ascent approach, aiming to approximately optimize the exact ranking measures, for large scale problems. There are also algorithms for *ordinal regression*, which is a related but different learning problem (Li et al, 2007; Crammer et al, 2001; Herbrich et al, 1999), and *listwise* approaches to ranking (Cao et al, 2007; Xia et al, 2008; Xu and Li, 2007; Yue et al, 2007).

2 Learning-to-Rank and Learning-To-Rerank

We first introduce *learning-to-rank*, or *supervised bipartite ranking*. The training data are labeled observations $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, with observations $\mathbf{x}_i \in X \subset \mathcal{R}^d$ and labels $y_i \in \{0, 1\}$ for all i . The observations labeled “1” are called “positive observations,” and the observations labeled “0” are “negative observations.” There are n_+ positive observations and n_- negative observations, with index sets $S_+ = \{i : y_i = 1\}$ and $S_- = \{k : y_k = 0\}$. A ranking algorithm uses the training data to produce a scoring function $f : X \rightarrow \mathcal{R}$ that assigns each observation a real-valued score. Ideally, for a set of test observations drawn from the same (unknown) distribution as the training data, f should rank the observations in order of $P(y = 1|\mathbf{x})$, and we measure the quality of the solution using “rank statistics,” or functions of the observations relative to each other. Note that bipartite ranking and binary classification are fundamentally different, and there are many works that explain the differences (e.g., Ertekin and Rudin, 2011). Briefly, classification algorithms consider a statistic of the observations relative to a decision boundary (n comparisons) whereas ranking algorithms consider observations relative to each other (on the order of n^2 comparisons for pairwise rank statistics).

Since the evaluation of test observations uses a chosen rank statistic, the same rank statistic (or a convexified version of it) is optimized on the training set to produce f . Regularization is added to help with generalization. Thus, a ranking algorithm looks like:

$$\min_{f \in \mathcal{F}} \text{RankStatistic}(f, \{\mathbf{x}_i, y_i\}_i) + C \cdot \text{Regularizer}(f).$$

This is the form of algorithm we consider for the reranking step.

2.1 Reranking

We are considering *reranking* methods, which have two ranking steps. In the first ranking step, a base algorithm is run over the training set, and a scoring function f_{initial} is produced and observations are rank-ordered by the score. A threshold is chosen, and all observations with scores above the threshold are reranked by another ranking algorithm which produces another scoring function f . To evaluate the quality of the solution on the test set, each test observation is evaluated first

by f_{initial} . For the observations with scores above the threshold, they are reranked according to f . The full ranking of test observations is produced by appending the test observations scored by f to the test observations scored only by f_{initial} .

2.2 Rank Statistics

We will extend the definition of conditional linear rank statistics (Clemençon and Vayatis, 2008; Clemençon et al, 2008) to include various definitions of rank. For now, we assume that there are no ties in score for any pair of observations, but we will heavily discuss ties later, and extend this definition to include rank definitions when there are ties. For the purpose of this section, the rank is currently defined so that the top of the list has the highest ranks, and all ranks are unique. The rank of an observation is the number of observations with scores at or beneath it:

$$\text{Rank}(f(\mathbf{x}_i)) = \sum_{t=1}^n \mathbf{1}_{[f(\mathbf{x}_t) \leq f(\mathbf{x}_i)]}.$$

Thus, ranks can range from 1 at the bottom to n at the top. A **conditional linear rank statistic** (CLRS) created from scoring function $f : X \rightarrow \mathcal{R}$ is of the form

$$\text{CLRS}(f) = \sum_{i=1}^n \mathbf{1}_{y_i=1} \phi(\text{Rank}(f(\mathbf{x}_i))).$$

Here ϕ is a non-decreasing function producing only non-negative values. Without loss of generality, we define $a_\ell := \phi(\ell)$, the contribution to the score if the observation with rank ℓ has label +1. By properties of ϕ , we know $0 \leq a_1 \leq a_2 \leq \dots \leq a_n$. Then

$$\text{CLRS}(f) = \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{Rank}(f(\mathbf{x}_i))=\ell]} \cdot a_\ell. \quad (1)$$

This class captures a broad collection of rank statistics, including the following well-known rank statistics:

- $a_\ell = \ell$: Wilcoxon Rank Sum (WRS) statistic, which is an affine function of the Area Under the Receiver Operator Characteristic Curve (AUC) when there are no ties in rank (that is, f such that $f(\mathbf{x}_i) \neq f(\mathbf{x}_k) \forall i \neq k$).

$$\text{WRS}(f) = \sum_{i \in S_+} \text{Rank}(f(\mathbf{x}_i)) = n_+ n_- \cdot \text{AUC}(f) + \frac{n_+(n_+ + 1)}{2}.$$

If ties are present, we would subtract the number of ties within the positive class from the right side of the equation above. The AUC is the fraction of correctly ranked positive-negative pairs:

$$\text{AUC}(f) = \frac{1}{n_+ n_-} \sum_{i \in S_+} \sum_{k \in S_-} \mathbf{1}_{[f(\mathbf{x}_k) < f(\mathbf{x}_i)]}.$$

The AUC, when multiplied by constant $n_+ n_-$, is the Mann-Whitney U statistic. The AUC has an affine relationship with the pairwise misranking error (the fraction of positive-negative pairs in which a positive is ranked at or below a negative):

$$\text{PairwiseMisrankingError}(f) = 1 - \text{AUC}(f) = \frac{1}{n_+ n_-} \sum_{i \in S_+} \sum_{k \in S_-} \mathbf{1}_{[f(\mathbf{x}_k) \geq f(\mathbf{x}_i)]}. \quad (2)$$

Some ranking algorithms are designed to approximately minimize the pairwise misranking error, e.g., RankBoost (Freund et al, 2003).

- $a_\ell = \ell \cdot \mathbf{1}_{[\ell \geq \theta]}$ for predetermined threshold θ : Related to the local AUC or partial AUC, which looks at the area under the leftmost part of the ROC curve only (Clemençon and Vayatis, 2007, 2008; Dodd and Pepe, 2003). The leftmost part of the ROC curve is the top portion of the ranked list. The top of the list is the most important in applications such as information retrieval and maintenance.
- $a_\ell = \mathbf{1}_{[\ell=n]}$: Winner Takes All (WTA), which is 1 when the top observation in the list is positively-labeled and 0 otherwise (Burges et al, 2006).
- $a_\ell = \frac{1}{n-\ell+1}$: Mean Reciprocal Rank (MRR) (Burges et al, 2006).
- $a_\ell = \frac{1}{\log_2(n-\ell+2)}$: Discounted Cumulative Gain (DCG), which is used in information retrieval (Järvelin and Kekäläinen, 2000).
- $a_\ell = \frac{1}{\log_2(n-\ell+2)} \cdot \mathbf{1}_{[\ell \geq N]}$: DCG@N, which cuts off the DCG after the top N. (See, for instance, Le et al, 2010).
- $a_\ell = \ell^p$ for some $p > 0$: Similar to the P -Norm Push, which uses ℓ_p norms to focus on the top of the list, the same way as an ℓ_p norm focuses on the largest elements of a vector (Rudin, 2009a).

Rank statistics have been studied in several theoretical papers (e.g., Wang et al, 2013).

2.3 Some Known Methods for Learning-To-Rank

Current methods for learning-to-rank optimize convex proxies for the rank statistics discussed above. RankBoost (Freund et al, 2003) uses the exponential loss function as an upper bound for the 0-1 loss within the misranking error, $\mathbf{1}_{[z \leq 0]} \leq e^{-z}$, and minimizes

$$\sum_{i \in S_+} \sum_{k \in S_-} e^{-(f(\mathbf{x}_i) - f(\mathbf{x}_k))}, \quad (3)$$

whereas support vector machine ranking algorithms (e.g., Joachims, 2002; Herbrich et al, 2000; Shen and Joshi, 2003) use the hinge loss $\max\{0, 1 - z\}$, that is:

$$\sum_{i \in S_+} \sum_{k \in S_-} \max\{0, 1 - (f(\mathbf{x}_i) - f(\mathbf{x}_k))\} + C \|f\|_K^2, \quad (4)$$

where the regularization term is a reproducing kernel Hilbert space norm. Other ranking algorithms include RankProp and RankNet (Caruana et al, 1996; Burges et al, 2005).

We note that the class of CLRS includes a very wide range of rank statistics, some of which concentrate on the top of the list (e.g., DCG) and some that do not (e.g., WRS), and it is not clear which conditional linear rank statistics (if any) from the CLRS are close to the convexified loss functions (3) and (4).

Since the convexified loss functions do not necessarily represent the rank statistics of interest, it is not even necessarily true that an algorithm for ranking will perform better for ranking than an algorithm designed for classification; in fact, AdaBoost and RankBoost provably perform equally well for ranking under fairly general circumstances (Rudin and Schapire, 2009). Ertekin and Rudin (2011) provide a discussion and comparison of classification versus ranking methods. Ranking algorithms ultimately aim to put the observations in order of $P(y = 1|x)$, and so do some classification algorithms such as logistic regression. Thus, one might consider using logistic regression for ranking (e.g., Cooper et al, 1994; Fine et al, 1997; Perlich et al, 2003). Logistic regression minimizes:

$$\sum_{i=1}^n \ln \left(1 + e^{-y_i f(\mathbf{x}_i)} \right). \quad (5)$$

This loss function does not closely resemble the AUC. On the other hand, it is surprising how common it is within the literature to use logistic regression to produce a predictive model, and yet evaluate the quality of the learned model using AUC.

Since RankBoost, RankProp, RankNet, etc., do not directly optimize any CLRS, they do not have the problem with ties in score that we will find when we directly try to optimize a CLRS.

2.4 Why Learning-To-Rank Methods Can Fail

We prove that the exponential loss and other common loss functions may yield poor results for some rank statistics.

Theorem 1 *There is a simple one-dimensional dataset for which there exist two ranked lists (called Solution 1 and Solution 2) that are completely reversed from each other (the top of one list is the bottom of the other and vice versa) such that the WRS (the AUC), partial AUC@100, DCG, MRR and hinge loss prefer Solution 1, whereas the DCG@100, partialAUC@10 and exponential loss all prefer Solution 2.*

The proof is by construction. Along the single dimension x , the dataset has 10 negatives near $x=3$, then 3000 positives near $x=1$, then 3000 negatives near $x=0$, and 80 positives near $x=-10$. We generated each of the four clumps of points with a standard deviation of 0.05 just so that there would not be ties in score. Figure 1 shows data drawn from the distribution, where for display purposes we spread the points along the horizontal axis, but the vertical axis is the only one that matters: one ranked list goes from top to bottom (Solution 1) and the other goes from bottom to top (Solution 2).

The bigger clumps are designed to dominate rank statistics that do not decay (or decay slowly) down the list, like the WRS. The smaller clumps are designed to dominate rank statistics that concentrate on the top of the list, like the partial WRS or partial DCG.

This theorem means that using the exponential loss to approximate the AUC, as RankBoost does, could give the completely opposite result than desired. It also means that using the hinge loss to approximate the partial DCG or partial AUC could yield completely the wrong result. Further, the fact that the exponential loss and hinge loss behave differently also suggests that convex losses can behave quite differently than the underlying rank statistics that they are meant to approximate.



Fig. 1 An illustrative distribution of data. Positive observations are gray and negative observations are black. Solution 1 ranks observations from top to bottom, and Solution 2 ranks solutions from bottom to top.

Another way to say this is that the convexification “washes out” the differences between rank statistics. If we were directly to optimize the rank statistic of interest, the problem discussed above would vanish.

It is not surprising that rank statistics can behave quite differently on the same dataset. Rank statistics are very different than classification statistics. Rank statistics consider every pair of observations relative to each other, so even small changes in a scoring function f can lead to large changes in a rank statistic. Classification is different – observations are considered relative only to a decision boundary.

The example considered in this section also illustrates why arguments about consistency (or lack thereof) of ranking methods (e.g., Kotlowski et al, 2011) are not generally relevant for practice. Sometimes these arguments rely on incorrect assumptions about the class of models used for ranking with respect to the underlying distribution of the data. These arguments also depend on how the modeler is assumed to “change” this class as the sample size increases to infinity. The tightest bounds available for limited function classes and for finite data are those from statistical learning theory. Those bounds support optimizing rank statistics.

To optimize rank statistics, there is a need for more refined models; however, this refinement comes at a computational cost of solving a harder problem. This thought has been considered in several previous works on learning-to-rank (Le et al, 2010; Ertekin and Rudin, 2011; Tan et al, 2013; Chakrabarti et al, 2008; Qin et al, 2013).

2.5 Most Learning-To-Rank Methods Have The Problem Discussed Above

The class of CLRS includes a very wide range of rank statistics, some of which concentrate on the top of the list (e.g., DCG) and some that do not (e.g., WRS), and it is not clear which conditional linear rank statistics (if any) from the CLRS are close to the convexified loss functions of the ranking algorithms. RankBoost is not the only algorithm where problems can occur, and they can also occur for support vector machine ranking algorithms (e.g., Joachims, 2002; Herbrich et al, 2000) and algorithms like RankProp and RankNet (Caruana et al, 1996; Burges et al, 2005). The methods of Ataman et al (2006), Brooks (2010), and Tan et al (2013) have used linear relaxations or greedy methods for learning to rank, rather than exact reranking, which will have similar issues; if one optimizes the wrong rank statistic, one may not achieve the correct answer. Logistic regression is commonly used for ranking. Logistic regression minimizes: $\sum_{i=1}^n \ln(1 + e^{-y_i f(\mathbf{x}_i)})$. This loss function does not closely resemble AUC. On the other hand, it is surprising how common it is to use logistic regression to produce a predictive model, and yet evaluate the quality of the model using AUC.

The fundamental premise of learning-to-rank is that better test performance can be achieved by optimizing the performance measure (a rank statistic) on the training set. This means that one should choose to optimize differently for each rank statistic. However, in practice when the same convex substitute is used to approximate a variety of rank statistics, it directly undermines this fundamental premise, and could compromise the quality of the solution. If convexified rank statistics are a reasonable substitute for rank statistics, we would expect to see that (i) the rank statistics are reasonably approximated by their convexified versions, (ii) if we consider several convex proxies for the same rank statistic (in this case AUC), then they should all behave very similarly to each other, and similarly to the true (non-convexified) AUC. However, as we discussed, neither of these are true.

2.6 Ties and Problematic, Thus Use ResolvedRank and Subrank

Dealing with ties in rank is critical when directly optimizing rank statistics. If a tie in rank between a positive and negative is considered as correct, then an optimal learning algorithm would produce the trivial scoring function $f(x) = \text{constant } \forall x$;

Label y_i	+	+	-	-	-	+	+	-	+
Score $f(\mathbf{x}_i)$	6.2	6.2	5.8	4.6	3.1	3.1	2.3	1.7	1.7
SubRank	7	7	6	5	3	3	2	0	0
ResolvedRank	8	7	6	5	4	3	2	1	0

Fig. 2 Demonstration of rank definitions.

this solution would unfortunately attain the highest possible score when optimizing any pairwise rank statistic. This problem happens, for instance, with the definition of Clemençon and Vayatis (2008), that is:

$$\text{RankCV}(f(\mathbf{x}_i)) = \sum_{k=1}^n \mathbf{1}_{f(\mathbf{x}_k) \leq f(\mathbf{x}_i)},$$

which counts ties in score as correct. Using this definition for rank in the CLRS:

$$\text{CLRS}_{\text{CV}}(f) = \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{RankCV}(f(\mathbf{x}_i))=\ell]} \cdot a_\ell. \quad (6)$$

we find that optimizing CLRS_{CV} directly yields the trivial solution that all observations get the same score. So this definition of rank should not be used.

We need to encourage our ranking algorithm not to produce ties in score, and thus in rank. To do this, we pessimistically consider a tie between our positive and a negative as a misrank. We will use two definitions of rank within the CLRS – *ResolvedRanks* and *Subranks*. For *ResolvedRanks*, when negatives are tied with positives, we force the negatives to be higher ranked. For *Subranks*, we do not force this, but when we optimize the CLRS, we will prove that ties are resolved this way anyway.

The assignment of *ResolvedRanks* and *Subranks* are not unique, there can be multiple ways to assign *ResolvedRanks* or *Subranks* for a set of observations.

We define the **Subrank** by the following formula:

$$\text{Subrank}(f(\mathbf{x}_i)) = \sum_{k=1}^n \mathbf{1}_{[f(\mathbf{x}_k) < f(\mathbf{x}_i)]}, \quad \forall i = 1, \dots, n.$$

The Subrank of observation i is the number of observations that score strictly below it. Subranks range from 0 to $n - 1$ and the CLRS becomes:

$$\text{CLRS}_{\text{Subrank}}(f) = \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{Subrank}(f(\mathbf{x}_i))=\ell-1]} \cdot a_\ell. \quad (7)$$

Observations with equal score have tied Subranks.

ResolvedRanks are defined as follows, where the tied ranks are resolved pessimistically. **ResolvedRanks** are assigned so that:

1. The *ResolvedRank* of an observation is greater than or equal to its *Subrank*.
2. If a positive observation and a negative observation have the same score, then the negative observation gets a higher *ResolvedRank*.
3. Each possible *ResolvedRank*, 0 through $n - 1$, is assigned to exactly one observation.

The SubRanks and ResolvedRanks are equal to each other when there are no ties in score. We provide one possible assignment of Subranks and ResolvedRanks in Figure 2 to demonstrate the treatment of ties. We then have the CLRS with ResolvedRanks as:

$$\text{CLRS}_{\text{ResolvedRank}}(f) = \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(f(\mathbf{x}_i))=\ell-1]} \cdot a_\ell. \quad (8)$$

The ResolvedRanks are the quantity of interest, as optimizing them will provide a scoring function with minimal misranks and minimal ties between positives and negatives.

Note that ties are not fundamental in other statistical uses of rank statistics, such as hypothesis testing. Ties are usually addressed by fixing them, or assigning the tied observations a (possibly fractional) rank that is the average (e.g., tied observations would get ranks 7.5 rather than 7 and 8) (see Tamhane and Dunlop, 2000; Wackerly et al, 2002). Ties are not treated uniformly across statistical applications (Savage, 1957), and there has been comparative work on treatment of ties (e.g., Putter, 1955). This differs from when we optimize rank statistics, where ties are of central importance as we discussed.

3 Reranking Formulations Using ResolvedRanks and Subranks

Here we produce the two formulations – one for optimizing the regularized CLRS with ResolvedRanks, and the other for optimizing the regularized CLRS with Subranks.

3.1 Maximize the Regularized CLRS with ResolvedRanks

We would like to optimize the general CLRS, for any choices of the a_ℓ 's, where we want to penalize ties in rank between positives and negatives, and we would also like a full ranking of observations. Thus, we will directly optimize $\text{CLRS}_{\text{ResolvedRank}}(f) + C \cdot \text{Regularizer}(f)$ for our reranking algorithm. Our hypothesis space is the space of linear scoring functions $f(\mathbf{x}_i) = \mathbf{w}^T \mathbf{x}_i$, where $\mathbf{w} \in \mathcal{R}^d$.

$$\begin{aligned} & \max_{\mathbf{w} \in \mathcal{R}^d} \text{CLRS}_{\text{ResolvedRank}}(\mathbf{w}) - C \|\mathbf{w}\|_0 \\ &= \max_{\mathbf{w} \in \mathcal{R}^d} \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(\mathbf{w}^T \mathbf{x}_i) = \ell - 1]} \cdot a_\ell - C \|\mathbf{w}\|_0. \end{aligned}$$

Nonlinearities can be incorporated as usual by including additional variables, such as indicator variables or nonlinear functions of the original variables. We optimize over choices for vector \mathbf{w} .

Building up to the formulation, we will create the binary variable $t_{i\ell}$ so that it is 1 for $\ell \leq \text{ResolvedRank}(f(\mathbf{x}_i)) + 1$ and 0 otherwise. That is, if observation i has ResolvedRank equal to 5, then t_{i1}, \dots, t_{i6} are all 1 and t_{i7}, \dots, t_{in} are 0. Then

$$\sum_{\ell=1}^n (a_\ell - a_{\ell-1}) t_{i\ell} \tag{9}$$

is a telescoping sum for $\ell \leq \text{ResolvedRank}(f(\mathbf{x}_i)) + 1$. When we define $a_0 = 0$, the sum (9) becomes simply $a_{\text{ResolvedRank}(f(\mathbf{x}_i)) + 1}$, or equivalently, the term from (8):

$$\sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(f(\mathbf{x}_i)) = \ell - 1]} \cdot a_\ell.$$

As in (8) we multiply by y_i and sum over observations to produce the $\text{CLRS}_{\text{ResolvedRank}}$. Doing this to (9), $\text{CLRS}_{\text{ResolvedRank}}$ becomes:

$$\sum_{i=1}^n y_i \sum_{\ell=1}^n (a_\ell - a_{\ell-1}) t_{i\ell} \quad \text{where } a_0 = 0.$$

By definition $t_{i1} = 1$ for all i , so we can simplify the $\text{CLRS}_{\text{ResolvedRank}}$ function above to:

$$\sum_{i \in S_+} \left(\sum_{\ell=2}^n (a_\ell - a_{\ell-1}) t_{i\ell} + a_1 \right) = |S_+| a_1 + \sum_{i \in S_+} \sum_{\ell=2}^n (a_\ell - a_{\ell-1}) t_{i\ell}.$$

Note that the differences $a_\ell - a_{\ell-1}$ are all nonnegative. When they are zero they cannot contribute to the $\text{CLRS}_{\text{ResolvedRank}}$ function. When they are strictly positive there can be a contribution made to the $\text{CLRS}_{\text{ResolvedRank}}$ function. Thus, we introduce notation $\tilde{a}_\ell = a_\ell - a_{\ell-1}$ and $S_r = \{\ell \geq 2 : \tilde{a}_\ell > 0\}$ which are used in both formulations below. The $\text{CLRS}_{\text{ResolvedRank}}$ becomes:

$$|S_+| a_1 + \sum_{i \in S_+} \sum_{\ell \in S_r} \tilde{a}_\ell t_{i\ell}. \tag{10}$$

We will maximize this, which means that the $t_{i\ell}$'s will be set to 1 when possible, because the \tilde{a}_ℓ 's in the sum are all positive. When we maximize, we do not need the constant $|S_+| a_1$ term.

We define integer variables $r_i \in [0, n - 1]$ to represent the ResolvedRanks of the observations. Variables r_i and $t_{i\ell}$ are related in that $t_{i\ell}$ can only be 1 when $\ell \leq r_i + 1$, implying $t_{i\ell} \leq \frac{r_i}{\ell - 1}$.

We use linear scoring functions, so the score of instance \mathbf{x}_i is $\mathbf{w}^T \mathbf{x}_i$. Variables z_{ik} are indicators of whether the score of observation i is above the score of observation k . Thus we want to have $z_{ik} = 1$ if $\mathbf{w}^T \mathbf{x}_i > \mathbf{w}^T \mathbf{x}_k$ and $z_{ik} = 0$ otherwise. Beyond this we want to ensure no ties in score, so we want all scores to be at least ε apart. This will be discussed further momentarily.

Our first ranking algorithm is below, which maximizes the regularized CLRS using ResolvedRanks.

$$\begin{aligned} \operatorname{argmax}_{\mathbf{w}, \gamma_j, z_{ik}, t_{i\ell}, r_i \forall i, k, \ell, j} \quad & \sum_{i \in S_+} \sum_{\ell \in S_r} \tilde{a}_i t_{i\ell} - C \sum_j \gamma_j \quad \text{subject to} & (11) \\ z_{ik} \leq \mathbf{w}^T (\mathbf{x}_i - \mathbf{x}_k) + 1 - \varepsilon, \quad & \forall i, k = 1, \dots, n, & (12) \\ z_{ik} \geq \mathbf{w}^T (\mathbf{x}_i - \mathbf{x}_k), \quad & \forall i, k = 1, \dots, n, & (13) \\ \gamma_j \geq w_j & & (14) \\ \gamma_j \geq -w_j & & (15) \\ r_i - r_k \geq 1 + n(z_{ik} - 1), \quad & \forall i, k = 1, \dots, n, & (16) \\ r_k - r_i \geq 1 - nz_{ik}, \quad & \forall i \in S_+, k \in S_-, & (17) \\ r_k - r_i \geq 1 - nz_{ik}, \quad & \forall i, k \in S_+, i < k, & (18) \\ r_k - r_i \geq 1 - nz_{ik}, \quad & \forall i, k \in S_-, i < k, & (19) \\ t_{i\ell} \leq \frac{r_i}{\ell - 1}, \quad & \forall i \in S_+, \ell \in S_r, & (20) \\ -1 \leq w_j \leq 1, \quad & \forall j = 1, \dots, d, \\ 0 \leq r_i \leq n - 1, \quad & \forall i = 1, \dots, n, \\ z_{ik} \in \{0, 1\}, \quad & \forall i, k = 1, \dots, n, \\ t_{i\ell} \in \{0, 1\}, \quad & \forall i \in S_+, \ell \in S_r, \\ \gamma_j \in \{0, 1\}, \quad & \forall j \in \{1, \dots, d\}. & (21) \end{aligned}$$

To ensure that solutions with ranks that are close together are not feasible, Constraint (12) forces $z_{ik} = 0$ if $\mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_k < \varepsilon$, and Constraint (13) forces $z_{ik} = 1$ if $\mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_k > 0$. Thus, a solution where any two observations have a score difference above 0 and less than ε is not feasible. (Note that these constraints alone do not prevent a score difference of exactly 0; for that we need the constraints that follow.) Constraints (14) and (15) define the γ_j 's to be indicators of nonzero coefficients w_j .

Constraints (16)-(19) are the ‘‘tie resolution’’ equations. Constraint (16) says that for any pair $(\mathbf{x}_i, \mathbf{x}_k)$, if the score of i is larger than that of k so that $z_{ik} = 1$, then $r_i \geq r_k + 1$. That handles the assignment of ranks when there are no ties, so now we need only to resolve ties in the score. We have Constraint (17) that applies to positive-negative pairs: when the pair is tied, this constraint forces the negative observation to have higher rank. Similarly, Constraints (18) and (19) apply to positive-positive pairs and negative-negative pairs respectively, and state that ties are broken lexicographically, that is, according to their index in the dataset.

We discussed Constraint (20) earlier, which provides the definition of $t_{i\ell}$ so that $t_{i\ell} = 1$ whenever $\ell \leq r_i + 1$. Also we force the w_j 's to be between -1 and 1 so their values do not go to infinity and so that the ε values are meaningful, in that they can be considered relative to the maximum possible range of w_j .

3.2 Maximize the Regularized CLRS with Subranks

We are solving:

$$\begin{aligned} & \max_{\mathbf{w} \in \mathcal{R}^d} \text{CLRS}_{\text{Subrank}}(\mathbf{w}) - C \|\mathbf{w}\|_0 \\ & = \max_{\mathbf{w} \in \mathcal{R}^d} \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{Subrank}(\mathbf{w}^T \mathbf{x}_i) = \ell - 1]} \cdot a_\ell - C \|\mathbf{w}\|_0. \end{aligned}$$

Maximizing the Subrank problem is much easier, since we do not want to force a unique assignment of ranks. This means the ‘‘tie resolution’’ equations are no longer present. We can directly assign a Subrank for observation i by $r_i = \sum_{k=1}^n z_{ik}$ because it is exactly the count of observations ranked beneath observation i ; that way the r_i variables do not even need to appear in the formulation.

Here is the formulation:

$$\operatorname{argmax}_{\mathbf{w}, \gamma_j, z_{ik}, t_{i\ell} \forall i, k, \ell, j} \sum_{i \in S_+} \sum_{\ell \in S_r} \tilde{a}_i t_{i\ell} - C \sum_j \gamma_j \quad \text{subject to} \quad (22)$$

$$t_{i\ell} \leq \frac{1}{\ell-1} \sum_{k=1}^n z_{ik}, \quad \forall i \in S_+, \ell \in S_r, \quad (23)$$

$$z_{ik} \leq \mathbf{w}^T (\mathbf{x}_i - \mathbf{x}_k) + 1 - \varepsilon, \quad \forall i \in S_+, k = 1, \dots, n, \quad (24)$$

$$\gamma_j \geq w_j \quad (25)$$

$$\gamma_j \geq -w_j \quad (26)$$

$$z_{ik} + z_{ki} = \mathbf{1}_{[\mathbf{x}_i \neq \mathbf{x}_k]}, \quad \forall i, k \in S_+, \quad (27)$$

$$t_{i\ell} \geq t_{i, \ell+1}, \quad \forall i \in S_+, \ell \in S_r \setminus \max(\ell \in S_r), \quad (28)$$

$$\sum_{i \in S_+} \sum_{\ell \in S_r} \tilde{a}_i t_{i\ell} \leq \sum_{\ell=1}^n a_\ell, \quad (29)$$

$$z_{ik} = 0, \quad \forall i \in S_+, k = 1, \dots, n, \mathbf{x}_i = \mathbf{x}_k, \quad (30)$$

$$-1 \leq w_j \leq 1, \quad \forall j = 1, \dots, d,$$

$$t_{i\ell}, z_{ik}, \gamma_j \in \{0, 1\}, \quad \forall i \in S_+, \ell \in S_r, k = 1, \dots, n, j \in \{1, \dots, d\}.$$

Constraint (23) is similar to Constraint (20) from the ResolvedRank formulation. Since we are maximizing with respect to the $t_{i\ell}$'s, the z_{ik} 's will naturally be maximized by Constraint (23). Thus we need to again force the z_{ik} 's down to 0 when $\mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_k < \varepsilon$, which is done via Constraint (24). Constraints (25) and (26) define the γ_j 's to be indicators of nonzero coefficients w_j . It is not necessary to include Constraints (27) through (30); they are there only to speed up computation, by helping to make the linear relaxation of the integer program closer to the set of feasible integer points. For the experiments in this paper they did not substantially speed up computation and we chose not to use them.

Beyond the formulations presented here, we have placed a formulation for optimizing the regularized AUC in the Appendix A.1 and another formulation for optimizing the general pairwise rank statistic that inspired RankBoost (Freund et al, 2003) in Appendix A.2.

4 Why Subranks Are Often Sufficient

The ResolvedRank formulation above has $2d + n^2 + n_+ |S_r| + n$ variables, which is the total number of w, γ, z, t , and r variables. The Subrank formulation on the other hand has only $2d + n_+ n + n_+ |S_r|$ variables, since we only have w, γ, z , and t . This difference of $n_- \cdot n + n$ variables can heavily influence the speed at which we are able to find a solution. We would ultimately like to get away with solving the Subrank problem rather than the ResolvedRank problem. This would allow us to scale up our reranking problem substantially. In this section we will show why this is generally possible.

Denote the objectives as follows, where we have $f(\mathbf{x}_i) = \mathbf{w}^T \mathbf{x}_i$.

$$G_{\text{RR}}(f) := \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(f(\mathbf{x}_i))=\ell-1]} \cdot a_\ell - C \|\mathbf{w}\|_0$$

$$G_{\text{Sub}}(f) := \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{Subrank}(f(\mathbf{x}_i))=\ell-1]} \cdot a_\ell - C \|\mathbf{w}\|_0.$$

In this section, we will ultimately prove that any maximizer of G_{Sub} also maximizes G_{RR} . This is true under a very general condition, which is that there are no exactly duplicated observations. The reason for this condition is not completely obvious. In the Subrank formulation, if two observations are exactly the same, they will always get the same score and Subrank - there is no mechanism to resolve ties and assign ranks. This causes problems when approximating the ResolvedRank with the Subrank. We remark however, that this should not be a problem in practice. First, we can check in advance whether any of our observations are exact copies of each other, so we know whether it is likely to be a problem. Second, if we do have duplicated observations, we can always slightly perturb the \mathbf{x} values of the duplicated observations so they are not identical. Third, we remark that if the data are chosen from a continuous distribution, with probability 1 the observations will all be distinct anyway. We have found that in practice the Subrank formulation does not have problems even when there are ties.

In the first part of the section, we consider whether there are maximizers of G_{RR} that have no ties in score, in other words, solutions \mathbf{w} where $f(\mathbf{x}_i) \neq f(\mathbf{x}_k)$ for any two observations i and k . Assuming such solutions exist, we then show

that any maximizer of G_{Sub} is also a maximizer of G_{RR} . This is the result within Theorem 2. In the second part of the section, we show that the assumption we made for Theorem 2 is always satisfied, assuming no duplicated observations. That is, a maximizer of G_{RR} with no ties in score exists. The outline within our technical report (Chang et al, 2011) follows a similar outline but does not include regularization.

The following lemma establishes basic facts about the two objectives:

Lemma 1 $G_{\text{Sub}}(f) \leq G_{\text{RR}}(f)$ for all f . Further, $G_{\text{Sub}}(f) = G_{\text{RR}}(f)$ for all f with no ties.

Proof Choose any function f . Since by definition $\text{Subrank}(f(x_i)) \leq \text{ResolvedRank}(f(x_i)) \forall i$, and since the a_ℓ are nondecreasing,

$$\begin{aligned} \sum_{\ell=1}^n \mathbf{1}_{[\text{Subrank}(f(x_i))=\ell-1]} \cdot a_\ell &= a_{(\text{Subrank}(f(x_i))+1)} \\ &\leq a_{(\text{ResolvedRank}(f(x_i))+1)} \\ &= \sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(f(x_i))=\ell-1]} \cdot a_\ell \quad \forall i. \end{aligned} \tag{31}$$

Multiplying both sides by y_i , summing over i and subtracting the regularization term from both sides yields $G_{\text{Sub}}(f) \leq G_{\text{RR}}(f)$. When no ties are present (that is, $f(x_i) \neq f(x_k) \forall i \neq k$), Subranks and ResolvedRanks are equal, and the inequality above becomes an equality, and in that case, $G_{\text{Sub}}(f) = G_{\text{RR}}(f)$.

This lemma will be used within the following theorem which says that maximizers of G_{Sub} are maximizers of G_{RR} .

Theorem 2 Assume that the set $\text{argmax}_f G_{\text{RR}}(f)$ contains at least one function \bar{f} having no ties in score. Then any f^* such that $f^* \in \text{argmax}_f G_{\text{Sub}}(f)$ also obeys $f^* \in \text{argmax}_f G_{\text{RR}}(f)$.

Proof Assume there exists $\bar{f} \in \text{argmax}_f G_{\text{RR}}(f)$ such that there are no ties in score. Since \bar{f} is a maximizer of G_{RR} and does not have ties, it is also a maximizer of G_{Sub} by Lemma 1:

$$G_{\text{Sub}}(\bar{f}) = G_{\text{RR}}(\bar{f}) = \max_f G_{\text{RR}}(f) \geq \max_f G_{\text{Sub}}(f), \text{ thus } G_{\text{Sub}}(\bar{f}) = \max_f G_{\text{Sub}}(f).$$

Let f^* be an arbitrary maximizer of $G_{\text{Sub}}(f)$ (not necessarily tie-free). We claim that f^* is also a maximizer of G_{RR} . Otherwise,

$$G_{\text{RR}}(f^*) < G_{\text{RR}}(\bar{f}) \stackrel{(a)}{=} G_{\text{Sub}}(\bar{f}) \stackrel{(b)}{=} G_{\text{Sub}}(f^*) \stackrel{(c)}{\leq} G_{\text{RR}}(f^*),$$

which is a contradiction. Equation (a) comes from Lemma 1 applied to \bar{f} . Equation (b) comes from the fact that both \bar{f} and f^* are maximizers of G_{Sub} . Inequality (c) comes from Lemma 1 applied to f^* .

Interestingly enough, it is true that if \bar{f} maximizes $G_{\text{RR}}(f)$ and it has no ties, then \bar{f} also maximizes $G_{\text{Sub}}(f)$. In particular,

$$\max_f G_{\text{Sub}}(f) \leq \max_f G_{\text{RR}}(f) \leq G_{\text{RR}}(\bar{f}) = G_{\text{Sub}}(\bar{f}).$$

Note that so far, the results about G_{RR} and G_{Sub} hold for functions from any arbitrary set; we did not need to have $f = \mathbf{w}^T \mathbf{x}$ in the preceding computations. In what follows we take advantage of the fact that f is a linear combination of features in order to perturb the function away from ties in score. With this method we will be able to achieve the same maximal value of G_{RR} but with no ties.

Define M to be the maximum absolute value of the features, so that for all i, j , we have $|x_{ij}| \leq M$.

Lemma 2 If we are given $\bar{f} \in \text{argmax}_f G_{\text{RR}}(f)$ that yields a scoring function $\bar{f}(\mathbf{x}) = \bar{\mathbf{w}}^T \mathbf{x}$ with ties, it is possible to construct a perturbed scoring function \hat{f} that:

- i preserves all pairwise orderings, $\bar{f}(\mathbf{x}_i) > \bar{f}(\mathbf{x}_k) \Rightarrow \hat{f}(\mathbf{x}_i) > \hat{f}(\mathbf{x}_k)$,
- ii has no ties, $\hat{f}(\mathbf{x}_i) \neq \hat{f}(\mathbf{x}_k)$ for all i, k .
- iii has $\|\hat{\mathbf{w}}\|_0 = \|\bar{\mathbf{w}}\|_0$.

This result holds whenever no observations are duplicates of each other, $\mathbf{x}_i \neq \mathbf{x}_k \forall i, k$.

Proof We will construct $\hat{f}(\mathbf{x}) = \hat{\mathbf{w}}^T \mathbf{x}$ using the following procedure:

Step 1 Find the nonzero indices of $\bar{\mathbf{w}}$: let $\bar{J} := \{j : \bar{w}_j \neq 0\}$. Choose a unit vector \mathbf{v} in $\mathcal{R}^{|\bar{J}|}$ uniformly at random. Construct vector $\mathbf{u} \in \mathcal{R}^d$ to be equivalent to \mathbf{v} for \mathbf{u} restricted to the dimensions J and 0 otherwise.

Step 2 Choose real number δ to be between 0 and η , where

$$\eta = \min \left\{ \frac{\text{margin}_{\bar{\mathbf{w}}}}{2M\sqrt{d}}, \min_{j \in \bar{J}} |w_j| \right\}$$

where in the above expression

$$\text{margin}_{\bar{\mathbf{w}}} = \min_{\{i,k:\bar{f}(\mathbf{x}_i) > \bar{f}(\mathbf{x}_k)\}} (\bar{f}(\mathbf{x}_i) - \bar{f}(\mathbf{x}_k)).$$

Step 3 Construct $\hat{\mathbf{w}}$ as follows: $\hat{\mathbf{w}} = \bar{\mathbf{w}} + \delta \mathbf{u}$.

With probability one, we will show that $\hat{f}(\mathbf{x}) = \hat{\mathbf{w}}^T \mathbf{x}$ preserves pairwise orderings of \bar{f} but with no ties.

We will prove each part of the lemma separately.

Proof of (i) We choose any two observations \mathbf{x}_i and \mathbf{x}_k where $\bar{f}(\mathbf{x}_i) > \bar{f}(\mathbf{x}_k)$, and we need to show that $\hat{f}(\mathbf{x}_i) > \hat{f}(\mathbf{x}_k)$.

$$\begin{aligned} \hat{f}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_k) &= (\bar{\mathbf{w}} + \delta \mathbf{u})^T (\mathbf{x}_i - \mathbf{x}_k) = \bar{\mathbf{w}}^T (\mathbf{x}_i - \mathbf{x}_k) + \delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \\ &= \bar{f}(\mathbf{x}_i) - \bar{f}(\mathbf{x}_k) + \delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \geq \text{margin}_{\bar{\mathbf{w}}} + \delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k). \end{aligned} \quad (32)$$

In order to bound the right hand side away from zero we will use that:

$$\|\mathbf{x}_i - \mathbf{x}_k\|_2 = \left(\sum_{j=1}^d (x_{ij} - x_{kj})^2 \right)^{1/2} \leq \left(\sum_{j=1}^d (2M)^2 \right)^{1/2} = 2M\sqrt{d}. \quad (33)$$

Now,

$$\left| \delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \right| \stackrel{(a)}{\leq} \delta \|\mathbf{u}\|_2 \|\mathbf{x}_i - \mathbf{x}_k\|_2 \stackrel{(b)}{\leq} \delta \cdot 2M\sqrt{d} \stackrel{(c)}{<} \frac{\text{margin}_{\bar{\mathbf{w}}}}{2M\sqrt{d}} \cdot 2M\sqrt{d} = \text{margin}_{\bar{\mathbf{w}}}.$$

Here, inequality (a) follows from the Cauchy-Schwarz inequality, (b) follows from (33) and that $\|\mathbf{u}\|_2 = 1$, and (c) follows from the bound on δ from Step 2 of the procedure for constructing \hat{f} above. Thus $\delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) > -\text{margin}_{\bar{\mathbf{w}}}$, which combined with (32) yields

$$\hat{f}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_k) \geq \text{margin}_{\bar{\mathbf{w}}} + \delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) > \text{margin}_{\bar{\mathbf{w}}} - \text{margin}_{\bar{\mathbf{w}}} = 0.$$

Proof of (ii) We show that \hat{f} has no ties $\hat{f}(\mathbf{x}_i) \neq \hat{f}(\mathbf{x}_k)$ for all i, k . This must be true with probability 1 over the choice of random vector \mathbf{u} .

Since we know that all pairwise inequalities are preserved, we need to ensure only that ties become untied through the perturbation \mathbf{u} . Thus, let us consider tied observations \mathbf{x}_i and \mathbf{x}_k , so $\bar{f}(\mathbf{x}_i) = \bar{f}(\mathbf{x}_k)$. We need to show that they become untied: we need to show $|\hat{f}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_k)| > 0$. Consider $|\hat{f}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_k)|$:

$$\begin{aligned} |\hat{f}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_k)| &= \left| (\bar{\mathbf{w}} + \delta \mathbf{u})^T (\mathbf{x}_i - \mathbf{x}_k) \right| = \left| \bar{\mathbf{w}}^T (\mathbf{x}_i - \mathbf{x}_k) + \delta \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \right| \\ &= |\delta| \left| \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \right|. \end{aligned}$$

We now use the key assumption that no two observations are duplicates – this implies that at least one entry of vector $\mathbf{x}_i - \mathbf{x}_k$ is nonzero. Further, since \mathbf{u} is a random vector, the probability that it is orthogonal to vector $\mathbf{x}_i - \mathbf{x}_k$ is zero. So, with probability one with respect to the choice of \mathbf{u} , we have $\left| \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \right| > 0$. From the expression above,

$$|\hat{f}(\mathbf{x}_i) - \hat{f}(\mathbf{x}_k)| = |\delta| \left| \mathbf{u}^T (\mathbf{x}_i - \mathbf{x}_k) \right| > 0.$$

Proof of (iii) By our definitions, $\hat{\mathbf{w}} = \bar{\mathbf{w}} + \delta \mathbf{u}$, $\delta \leq \min_{j \in \bar{J}} |w_j|$, and \mathbf{u} is only nonzero in the components where $\bar{\mathbf{w}}$ is not 0. Each component of \mathbf{u} is nonzero with probability 1. For component j where $\bar{w}_j \neq 0$, we have $|\delta u_j| \leq \delta \|\mathbf{u}\|_2 \leq \delta \leq \min_{j' \in \bar{J}} \bar{w}_{j'} \leq \bar{w}_j$ which means $|\hat{w}_j| = |\bar{w}_j + \delta u_j| > 0$. So, for all components where $\bar{\mathbf{w}}$ is nonzero, we also have $\hat{\mathbf{w}}$ nonzero in those components. Further, for all components where $\bar{\mathbf{w}}$ is zero, we also have $\hat{\mathbf{w}}$ zero in those components. Thus $\|\bar{\mathbf{w}}\|_0 = \|\hat{\mathbf{w}}\|_0$.

The result below establishes the main result of the section, which is that if we optimize G_{Sub} , we get an optimizer of G_{RR} even though it is a much more complex optimization problem to optimize G_{RR} directly.

Theorem 3 Given $f^* \in \operatorname{argmax}_f G_{\text{Sub}}(f)$, then $f^* \in \operatorname{argmax}_f G_{\text{RR}}(f)$.
This holds when there are no duplicated observations, $x_i \neq x_k \forall i, k$ where $i \neq k$.

Proof We will show that the assumption of Theorem 2, which says that G_{RR} has a maximizer with no ties, is always true. This will give us the desired result. Let $\bar{f} \in \operatorname{argmax}_f G_{\text{RR}}(f)$. Either \bar{f} has no ties already, in which case there is nothing to prove, or it does have ties. If so, we can take its vector $\bar{\mathbf{w}}$ and perturb it using Lemma 2. The resulting vector $\hat{\mathbf{w}}$ has no ties. We need only to show that $\hat{\mathbf{w}}$ also maximizes G_{RR} . To do this we will show $G_{\text{RR}}(\hat{f}) \geq G_{\text{RR}}(\bar{f})$.

We know that

$$\begin{aligned} G_{\text{RR}}(\bar{f}) &= \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(\bar{f}(\mathbf{x}_i))=\ell-1]} \cdot a_\ell - c \|\bar{\mathbf{w}}\|_0 \\ &= \sum_{i \in S_+} a_{(\text{ResolvedRank}(\bar{f}(\mathbf{x}_i))+1)} - c \|\bar{\mathbf{w}}\|_0, \\ G_{\text{RR}}(\hat{f}) &= \sum_{i=1}^n y_i \sum_{\ell=1}^n \mathbf{1}_{[\text{ResolvedRank}(\hat{f}(\mathbf{x}_i))=\ell-1]} \cdot a_\ell - c \|\hat{\mathbf{w}}\|_0 \\ &= \sum_{i \in S_+} a_{(\text{ResolvedRank}(\hat{f}(\mathbf{x}_i))+1)} - c \|\hat{\mathbf{w}}\|_0, \end{aligned}$$

and $\|\bar{\mathbf{w}}\|_0 = \|\hat{\mathbf{w}}\|_0$ by Lemma 2. We know that $a_1 \leq a_2 \leq \dots \leq a_n$. Thus, as long as the ResolvedRanks of the positive observations according to \hat{f} are the same or higher than their ResolvedRanks according to \bar{f} , we are done.

Consider the untied observations of \bar{f} , which are $\{i : \bar{f}(\mathbf{x}_i) \neq \bar{f}(\mathbf{x}_k) \text{ for any } k\}$. Those observations have $\text{ResolvedRank}(\bar{f}(\mathbf{x}_i)) = \text{ResolvedRank}(\hat{f}(\mathbf{x}_i))$ by Lemma 2(i) which says that all pairwise orderings are preserved.

What remains is to consider the tied observations of \bar{f} , which are $\{i : \bar{f}(\mathbf{x}_i) = \bar{f}(\mathbf{x}_k) \text{ for some } k\}$. Consider a set of tied observations $\mathbf{x}_\alpha, \dots, \mathbf{x}_\zeta$ where $f(\mathbf{x}_\alpha) = \dots = f(\mathbf{x}_\zeta)$. If their labels are all equal, $y_\alpha = \dots = y_\zeta$, then regardless of how they are permuted to create the ResolvedRank in either \bar{f} or \hat{f} , the total contribution of those observations to the G_{RR} will be the same. If the labels in the set differ, then \bar{f} assigns ResolvedRanks pessimistically, so that the negatives all have ResolvedRanks above the positive (according to the definition of ResolvedRanks). This means that by perturbing the solution, \hat{f} could potentially increase the ranks of some of these tied positive observations. In that case, some of the a_ℓ 's of \hat{f} become larger than those of \bar{f} . Thus, $G_{\text{RR}}(\hat{f}) \geq G_{\text{RR}}(\bar{f})$ and we are done.

The result in Theorem 3 shows why optimizing G_{Sub} is sufficient to obtain the maximizer of G_{RR} . This provides the underpinning for use of the Subrank formulation.

5 Empirical Discussion of Learning-To-Rank

Through our experiments with the Subrank formulation, we made several observations, which we will present empirical results to support below.

Observation 1: There are some datasets where reranking can substantially improve the quality of the solution.

We present comparative results on the performance of several baseline ranking methods, namely Logistic Regression (LR), Support Vector Machines (SVM), RankBoost (RB), and the P-Norm Push for $p = 2$ and for the Subrank MIP formulations at 4 different levels of the cutoff K for reranking. For the SVM, we tried regularization parameters 10^{-1} , 10^{-2} , \dots , 10^{-6} and reported the best result. We chose datasets with the right level of imbalance so that not all of the top observations belong to a single class; this ensures that the rank statistics are meaningful at the top of the list. We used several datasets that are suitable for the type of method we are proposing, namely:

- ROC Flexibility: This dataset is designed specifically to show differences in rank statistics (Rudin, 2009b). Note that this dataset has ties, but the ties do not seem to influence the quality of the solution. (It is generally possible in practice to use the Subrank formulation even in the case of ties.) ($n = 500$, $d = 5$)
- Abalone19: This dataset is an imbalanced version of the Abalone dataset where the positive observations belong to class 19. It is available from the KEEL repository (Alcalá-Fdez et al, 2011). It contains information about sex, length, height, and weight, and the goal is to determine the age of the abalone (19). ($n = 4174$, $d = 8$)
- UIS from the UMass Aids Research Unit (Hosmer et al, 2013): This dataset contains information about each patient's age, race, depression level at admission, drug usage, number of prior drug treatments, and current treatment, and the label represents whether the patient remained drug free for 12 months afterwards. ($n = 575$, $d = 8$)

- Travel: This dataset is from a survey of transportation uses between three Australian cities (Hosmer et al, 2013). It contains information about what modes of traffic are used (e.g., public bus, airplane, train, car) which is what we aim to predict, and features include the travel time, waiting time at the terminal, the cost of transportation, the commuters’ household income level, and the size of the party involved in the commute. ($n = 840, d = 7$)
- NHANES (physical activity): This dataset contains health information about patients including physical activity levels, height, weight, age, gender, blood pressure, marital status, cholesterol, etc. (Hosmer et al, 2013). The goal is to predict whether the person is considered to be obese. ($n = 600, d = 21$)
- Pima Indians Diabetes, from the National Institute of Diabetes and Digestive and Kidney Diseases, available from the UCI Machine Learning Repository (Bache and Lichman, 2013): The goal is to predict whether a woman will test positive for diabetes during her pregnancy, based on measurements of her blood glucose concentration in an oral glucose tolerance test, her blood pressure, body mass index, age, and other characteristics. ($n = 768, d = 8$)
- Gaussians: This is a synthetic 2 dimensional dataset, with 1250 points subsampled from a population containing two big clumps of training examples, each entry of each observation drawn from a normal distribution with variance 0.5, where the positive clump (3000 points) was generated with mean (0,1), and the negative clump (3000 points) was generated with mean (0,0). These bigger clumps are designed to dominate the WRS. In addition, there is a smaller 10 point negative clump generated with mean (10,1) and noise components each drawn from a normal with standard deviation 0.05, and a positive clump of 200 points generated with mean (0,-3) and noise drawn with standard deviation 0.05. Note that we do not expect the “flipping” to occur here as it did in Section 2.4 since we are using DCG, which is much more difficult to distinguish from WRS than a steeper rank statistic. ($n = 1250, d = 2$)

For the MIP-based methods, we used logistic regression as the base ranker, and the reranker was learned from the top K . We varied K between 50, 100, 150, and we also used the full list. An exception is made for the Abalone19 data set, for which K varies between 250, 500 and 750 instead because Abalone19 is a highly imbalanced data set. We stopped the computation after 2 hours for each trial (1 hour for the ROC flexibility dataset), which gives a higher chance for the lower- K rerankers to solve to optimality. Most of the $K=50$ experiments for the ROC flexibility dataset solved to optimality within 5 minutes. The reported means and standard deviations were computed over 10 randomly chosen training and test splits, where the same splits were used for all datasets. We chose to evaluate according to the DCG measure as it is used heavily in information retrieval applications (Järvelin and Kekäläinen, 2000). Al-Maskari et al (2007) report that DCG is similar to the way humans evaluate search results, based on a study of user satisfaction. We used $C = 10^{-3}$ for the ROC Flexibility dataset, and $C = 10^{-4}$ for the other datasets. Note that for the DCG measure in particular, it is difficult to see a large improvement; for instance even on the extreme experiment in Section 2.4 the improvement in DCG from flipping the classifier completely upside down was only 16%.

Table 1 shows the results of our experiments, where we highlighted the best algorithm for each dataset on both training and test in bold, and used italics to represent test set results that are not statistically significantly worse than the best algorithm according to a matched pairs t-test with significance level $\alpha = 0.05$. In terms of predictive performance, the smaller K models performed consistently well on these data, achieving the best test performance on all of these datasets. On some of the datasets, we see a $\sim 10\%$ average performance improvement from reranking. (The magnitude is not too much different as from the experiment in Section 2.4 where the classifier flips upside down.) On the Travel dataset in particular, the $K=50$ reranking model had superior results over all of the baselines uniformly across all 10 trials.

The work of Chang et al (2012) shows the benefits of carrying the computation to optimality on a specialized application of MIP learning-to-rank for reverse-engineering product quality rankings.

Observation 2: There is a tradeoff between computation and quality of solution.

If the number of elements to rerank (denoted by K) is too small, the solution will not generalize as well. Theoretical results of Rudin (2009a) suggest that there is a tradeoff between how well we can generalize and how much the rank statistic is focused on the top of the ranked list. The main result of that work shows that if the rank statistic concentrates very much at the top of the list (like, for instance, the mean reciprocal rank statistic) then we require more observations in order to generalize well. If the number of observations is too small, learning-to-rank methods may not be beneficial over traditional learning methods like logistic regression. Further, if the number of observations is too small, then the variation from training to test will be much larger than the gain in training error from using the correct rank statistic; again in that case, learning-to-rank would not be beneficial.

If the number of elements K is too large, we will not be able to sufficiently solve the reranking problem within the allotted time, and the solution again could suffer. This reinforces our point that we should not refrain from solving hard problems, particularly on the scale of reranking, but certain hard problems are harder than others and the computation needs to be done carefully.

Table 1 Datasets for which reranking can make a difference

Dataset	Baseline methods				MIP-based methods				
	LR	SVM	RB	P-norm Push	K = 50	K = 100	K = 150		
	Full List	Full List	Full List	Full List	Full List	Full List	Full List		
ROC	train	31.21 ± 1.65	30.94 ± 1.57	29.00 ± 1.39	31.33 ± 31.43	31.96 ± 1.32	31.84 ± 1.36	31.65 ± 1.12	28.43 ± 1.80
	test	31.35 ± 1.48	31.10 ± 1.63	29.57 ± 1.43	31.43 ± 1.61	32.16 ± 1.31	32.09 ± 1.31	31.74 ± 1.65	28.96 ± 2.40
Abalone19 ^a	train	3.63 ± 0.43	3.41 ± 0.47	3.40 ± 0.65	3.44 ± 0.47	4.89 ± 0.58	4.45 ± 0.50	4.13 ± 0.53	2.54 ± 0.35
	test	2.96 ± 0.42	3.02 ± 0.53	2.66 ± 0.49	3.03 ± 0.50	3.08 ± 0.49	2.89 ± 0.35	2.76 ± 0.38	2.42 ± 0.48
UIS	train	18.86 ± 1.32	18.46 ± 1.38	19.44 ± 1.44	18.78 ± 1.40	20.45 ± 1.23	19.76 ± 1.27	19.26 ± 1.05	18.84 ± 1.44
	test	17.88 ± 1.11	17.81 ± 1.21	17.70 ± 1.40	17.89 ± 1.13	18.00 ± 1.31	18.64 ± 1.51	17.79 ± 1.73	17.89 ± 0.67
Travel	train	28.16 ± 1.60	27.59 ± 1.61	26.57 ± 1.60	28.09 ± 1.62	28.30 ± 1.63	28.24 ± 1.56	27.12 ± 1.45	26.94 ± 1.36
	test	27.32 ± 1.70	26.81 ± 1.76	24.95 ± 1.63	27.24 ± 1.66	27.61 ± 1.70	27.39 ± 1.60	26.00 ± 2.26	26.31 ± 1.83
NHANES	train	14.69 ± 1.63	13.83 ± 1.87	13.75 ± 2.01	14.46 ± 1.57	15.48 ± 1.69	15.02 ± 1.93	14.73 ± 1.59	13.87 ± 1.25
	test	13.06 ± 1.74	12.98 ± 1.79	12.10 ± 1.75	13.18 ± 1.82	13.26 ± 1.50	12.71 ± 1.61	12.94 ± 1.55	13.09 ± 1.82
Pima	train	35.50 ± 1.66	35.30 ± 1.61	35.80 ± 1.44	35.64 ± 1.67	35.75 ± 1.67	35.43 ± 1.69	34.85 ± 1.96	34.77 ± 2.03
	test	34.18 ± 1.81	34.03 ± 1.83	33.83 ± 1.65	34.24 ± 1.82	34.44 ± 1.76	33.56 ± 1.87	33.72 ± 2.21	33.65 ± 2.01
Gaussians	train	69.25 ± 2.70	69.28 ± 2.70	71.31 ± 2.15	69.24 ± 2.70	71.71 ± 2.22	71.76 ± 2.18	71.58 ± 2.29	64.70 ± 2.53
	test	64.69 ± 2.45	64.73 ± 2.45	67.13 ± 2.06	64.65 ± 2.43	68.03 ± 2.27	67.91 ± 2.27	67.79 ± 2.30	59.89 ± 2.26

^a We use $K = 250$, $K = 500$ and $K = 750$ for this data set because it is highly imbalanced.

Table 2 Datasets for which reranking does not make a difference

Dataset	Baseline methods				MIP-based methods			
	LR	SVM	RB	P-norm Push	K = 50	K = 100	K = 150	
	Full List	Full List	Full List	Full List	Full List	Full List	Full List	
Haberman	train	12.94 ± 1.07	12.95 ± 1.06	13.95 ± 1.24	12.94 ± 1.06	13.10 ± 1.19	13.02 ± 1.20	13.13 ± 0.92
	test	12.82 ± 1.15	12.63 ± 1.15	12.01 ± 1.06	12.64 ± 1.09	12.64 ± 1.47	12.80 ± 1.13	12.46 ± 1.05
Polypharm	train	19.43 ± 1.42	19.24 ± 1.53	19.11 ± 1.55	18.73 ± 1.25	19.21 ± 0.93	18.63 ± 1.25	18.04 ± 1.04
	test	17.23 ± 1.63	17.71 ± 1.56	17.40 ± 1.70	18.05 ± 1.33	17.37 ± 1.33	17.59 ± 1.25	17.16 ± 1.99
Glow500	train	17.26 ± 1.16	16.52 ± 1.41	17.23 ± 1.25	17.02 ± 1.44	18.20 ± 1.24	17.70 ± 1.22	17.08 ± 1.40
	test	16.68 ± 1.18	17.26 ± 1.15	17.01 ± 0.69	17.24 ± 1.27	16.79 ± 1.09	17.10 ± 1.66	16.78 ± 1.33

^a There are only 153 observations in the training data for the Haberman survival dataset. We did not do $K = 150$ because at that point it makes sense to run the MIP on the full dataset.

Again consider Table 1. Note that the $K = 50$ and $K = 100$ rerankers perform consistently well on these datasets, both in training and in testing. However, if K is set too large, the optimization on the training set will not be able to be solved close to optimality in the allotted time, and the quality of the solution will be degraded. This is an explicit tradeoff between computation and the quality of the solution.

Observation 3: There are some datasets for which the variance of the result is larger than the differences in the rank statistics themselves.

These are cases where better relative training values do not necessarily lead to better relative test values. In these cases we do not think it is worthwhile to use ranking algorithms at all, let alone *reranking* algorithms. For these datasets, logistic regression may suffice. The cases where reranking/ranking makes a difference are cases where the variance of the training and test values are small enough that we can reliably distinguish between the different rank statistics.

We present results on three datasets in Table 2, computed in the same way as the results in Table 1, for which various things have gone wrong, such as the optimizer not being able to achieve the best result on the training set, but even worse, the results are inconsistent between training and test. The algorithm that optimizes best over the training set is not the same algorithm that achieves the best out-of-sample quality. These are cases where the algorithms do not generalize well enough so that a ranking algorithm is needed. The datasets used here are the Haberman survival dataset from the UCI Machine Learning Repository (Bache and Lichman, 2013) ($n = 300$, $d = 3$), Poly-pharmacy study on drug consumption (Hosmer et al, 2013) ($n = 500$, $d = 13$), and data from the GLOW study on fracture risk (Hosmer et al, 2013) ($n = 500$, $d = 14$).

Observation 4: As long as the margin parameter ε is sufficiently small without being too small so that the solver will not recognize it, the quality of the solution is maintained. The regularization parameter C also can have an influence on the quality of the solution, and it is useful not to over-regularize or under-regularize.

Note that if ε is too large, the solver will not be able to force all of the inequalities to be strictly satisfied with margin ε . This could force many good solutions to be considered infeasible and this may ruin the quality of the solution. It could also cause problems with convergence of the optimization problem. When ε is smaller, it increases the size of the feasible solution space, so the problem is easier to solve. On the other hand, if ε is too small, the solver will have trouble recognizing the inequality and may have numerical problems.

In Table 3 we show what happens when the value of ε is varied on two of our datasets. We can see from Table 3 that as ε decreases by orders of magnitude the solution generally improves, but then at some point degrades. For the ROC Flexibility data, the $\varepsilon = 10^{-5}$ setting consistently performed better than the $\varepsilon = 10^{-6}$ setting over all 10 trials in both training and test. A similar observation holds for UIS, in that the $\varepsilon = 10^{-5}$ setting was able to optimize better than the $\varepsilon = 10^{-6}$ setting over all 10 trials on the training set.

Table 3 Different selections of ε


Dataset		10^{-1}	10^{-2}	10^{-3}	10^{-4}	10^{-5}	10^{-6}
ROC ¹	train	31.62 ± 1.25	31.85 ± 1.26	31.93 ± 1.36	31.84 ± 1.36	32.02 ± 1.30	31.58 ± 1.26
	test	31.59 ± 2.07	31.91 ± 1.57	32.10 ± 1.28	32.09 ± 1.31	32.21 ± 1.32	31.74 ± 1.26
UIS	train	19.70 ± 1.23	19.73 ± 1.24	19.80 ± 1.09	19.76 ± 1.27	19.73 ± 1.17	19.08 ± 1.40
	test	18.40 ± 1.09	18.03 ± 1.06	18.34 ± 1.20	18.64 ± 1.51	17.88 ± 1.66	18.23 ± 0.75

Table 4 shows the training and test performance as the regularization parameter C is varied over several orders of magnitude. As one would expect, a small amount of regularization helps performance, but too much regularization hurts performance as we start to sacrifice prediction quality for sparseness.

Table 4 Training and test performance for varying values of regularization parameter C .

Dataset		$C = 10^{-1}$	$C = 10^{-2}$	$C = 10^{-3}$	$C = 10^{-4}$	$C = 10^{-5}$	$C = 10^{-6}$
ROC	train	31.31 ± 1.52	31.31 ± 1.72	31.84 ± 1.36	31.94 ± 1.20	32.02 ± 1.30	31.86 ± 1.35
	test	31.35 ± 1.48	31.30 ± 1.57	32.09 ± 1.31	32.06 ± 1.53	32.21 ± 1.32	32.01 ± 1.35
UIS	train	19.15 ± 1.24	19.36 ± 1.02	19.69 ± 1.01	19.76 ± 1.27	19.89 ± 1.24	19.68 ± 1.08
	test	17.94 ± 1.11	18.15 ± 1.54	17.92 ± 1.57	18.64 ± 1.51	17.94 ± 1.12	17.74 ± 1.76

¹ The regularization constant C is set to 10^{-3} for this dataset.

Fig. 3 Objective values and optimality gap over time for ROC Flexibility dataset

Observation 5: Proving optimality takes longer than finding a reasonable solution.

Figure 3 shows the objective values and the upper bound on the optimality gap over time for four folds of the ROC Flexibility dataset, where K is 100 and C is 10^{-4} . Figure 4 shows the analogous plots for the UIS dataset. Usually a good solution is found within a few minutes, whereas proving optimality of the solution takes much longer. We do not require a proof of optimality to use the solution.

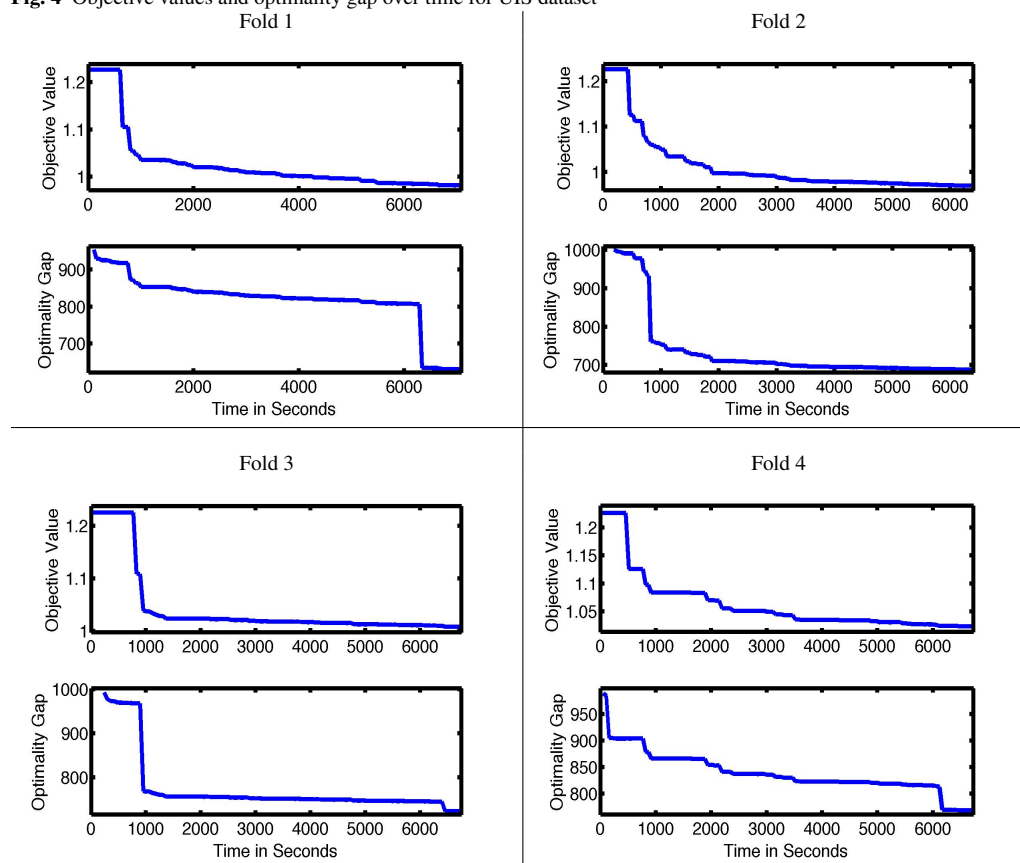
6 Conclusion

As shown through our discussion, using a computationally expensive reranking step may help to improve the quality of the solution for reranking problems. This can be useful in application domains such as maintenance prioritization and drug discovery where the extra time spent in obtaining the best possible solution can be very worthwhile. We proved an analytical reduction from the problem that we really want to solve (the ResolvedRank formulation) to a much more computationally tractable problem (the Subrank formulation). Through an experimental discussion, we explicitly showed the tradeoff between computation and the quality of the solution.

Acknowledgements

We would like to thank Allison Chang, our co-author on a technical report that inspired this work. This research was supported by the National Science Foundation under Grant No IIS-1053407 to C. Rudin, and an undergraduate exchange fellowship for Y. Wang through the Tsinghua-MIT-CUHK Research Center for Theoretical Computer Science.

Fig. 4 Objective values and optimality gap over time for UIS dataset



References

- Al-Maskari A, Sanderson M, Clough P (2007) The relationship between IR effectiveness measures and user satisfaction. In: Proceedings of SIGIR '07: Proceedings of the 30th International ACM SIGIR Conference on Research and Development in Information Retrieval, ACM, pp 773–774
- Alcalá-Fdez J, Fernández A, Luengo J, J Derrac SG, Sánchez L, Herrera F (2011) KEEL data-mining software tool: Data set repository, integration of algorithms and experimental analysis framework. *Journal of Multiple-Valued Logic and Soft Computing* 17(2–3):255–287
- Ataman K, Street WN, Zhang Y (2006) Learning to rank by maximizing AUC with linear programming. In: Proc. International Joint Conference on Neural Networks (IEEE IJCNN)
- Bache K, Lichman M (2013) UCI machine learning repository. URL <http://archive.ics.uci.edu/ml>
- Bradley AP (1997) The use of the area under the ROC curve in the evaluation of machine learning algorithms. *Pattern Recognition* 30(7):1145–1159
- Brooks JP (2010) Support vector machines with the ramp loss and the hard margin loss. *Operations Research* 59(2):467 – 479
- Burges C, Shaked T, Renshaw E, Lazier A, Deeds M, Hamilton N, Hullender G (2005) Learning to rank using gradient descent. In: Proc. 22nd International Conference on Machine Learning (ICML)
- Burges CJ, Ragnó R, Le QV (2006) Learning to rank with nonsmooth cost functions. In: Proc. Advances in Neural Information Processing Systems (NIPS), pp 395–402
- Cao Z, Qin T, Liu TY, Tsai MF, Li H (2007) Learning to rank: from pairwise approach to listwise approach. In: Proc. 24th International Conference on Machine Learning (ICML), ACM, New York, NY, USA, pp 129–136
- Caruana R, Baluja S, Mitchell T (1996) Using the future to “sort out” the present: Rankprop and multitask learning for medical risk evaluation. In: Advances in Neural Information Processing Systems (NIPS), vol 8, pp 959–965
- Chakrabarti S, Khanna R, Sawant U, Bhattacharyya C (2008) Structured learning for non-smooth ranking losses. In: Proc. 14th ACM SIGKDD international conference on Knowledge discovery and data mining (KDD), pp 88–96

- Chang A, Rudin C, Bertsimas D (2010) A discrete optimization approach to supervised ranking. In: Proc. INFORMS 5th Annual Workshop on Data Mining and Health Informatics
- Chang A, Rudin C, Bertsimas D (2011) Integer optimization methods for supervised ranking. Operations Research Center Working Paper Series OR 388-11, MIT
- Chang A, Rudin C, Cavaretta M, Thomas R, Chou G (2012) How to reverse-engineer quality rankings. *Machine Learning* 88:369–398
- Clemençon S, Vayatis N (2007) Ranking the best instances. *Journal of Machine Learning Research* 8:2671–2699
- Clemençon S, Vayatis N (2008) Empirical performance maximization for linear rank statistics. *Proc Advances in Neural Information Processing Systems (NIPS)* 21
- Clemençon S, Lugosi G, Vayatis N (2008) Ranking and empirical minimization of u-statistics. *Annals of Statistics* 36(2):844–874
- Collins M, Koo T (2005) Discriminative reranking for natural language parsing. *Journal of Association for Computational Linguistics* 31(1):25–70
- Cooper WS, Chen A, Gey FC (1994) Full text retrieval based on probabilistic equations with coefficients fitted by logistic regression. In: Proc. 2nd Text Retrieval Conference (TREC-2), pp 57–66
- Crammer K, Singer Y, et al (2001) Pranking with ranking. In: *Advances in Neural Information Processing Systems (NIPS)*, vol 1, pp 641–647
- Dodd LE, Pepe MS (2003) Partial AUC estimation and regression. UW Biostatistics Working Paper Series URL <http://www.bepress.com/uwbiostat/paper181>
- Ertekin Ş, Rudin C (2011) On equivalence relationships between classification and ranking algorithms. *Journal of Machine Learning Research* 12:2905–2929
- Fine MJ, Auble TE, Yealy DM, Hanusa BH, Weissfeld LA, Singer DE, Coley CM, Marrie TJ, Kapoor WN (1997) A prediction rule to identify low-risk patients with community-acquired pneumonia. *The New England Journal of Medicine* pp 243–250
- Freund Y, Iyer R, Schapire RE, Singer Y (2003) An efficient boosting algorithm for combining preferences. *Journal of Machine Learning Research* 4:933–969
- Herbrich R, Graepel T, Obermayer K (1999) Large margin rank boundaries for ordinal regression. *NIPS* pp 115–132
- Herbrich R, Graepel T, Obermayer K (2000) Large margin rank boundaries for ordinal regression. *Advances in Large Margin Classifiers*
- Hosmer DW, Lemeshow S, Sturdivant RX (2013) *Applied Logistic Regression: Third Edition*. John Wiley & Sons Inc.
- Jain V, Varma M (2011) Learning to re-rank: Query-dependent image re-ranking using click data. In: Proc. 20th International Conference on World Wide Web (WWW), pp 277–286
- Järvelin K, Kekäläinen J (2000) IR evaluation methods for retrieving highly relevant documents. In: Proc. 23rd Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp 41–48
- Ji H, Rudin C, Grishman R (2006) Re-ranking algorithms for name tagging. In: Proc. HLT-NAACL Workshop on Computationally Hard Problems and Joint Inference in Speech and Language Processing, pp 49–56
- Joachims T (2002) Optimizing search engines using clickthrough data. In: Proc. Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)
- Kang C, Wang X, Chen J, Liao C, Chang Y, Tseng B, Zheng Z (2011) Learning to re-rank web search results with multiple pairwise features. In: Proc. Fourth International Conference on Web Search and Web Data Mining (WSDM)
- Kessler JS, Nicolov N (2009) Targeting sentiment expressions through supervised ranking of linguistic configurations. In: *Proceedings of the Third International Conference on Weblogs and Social Media (ICWSM)*
- Kong D, Saar-Tsechansky M (2013) Collaborative information acquisition for data-driven decisions. *Machine Learning, Special Issue on ML for Science and Society*
- Kotlowski W, Dembczynski K, Hullermeier E (2011) Bipartite ranking through minimization of univariate loss. In: Proc. International Conference on Machine Learning
- Lafferty J, Zhai C (2001) Document language models, query models, and risk minimization for information retrieval. In: Proc. 24th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp 111–119
- Le QV, Smola A, Chapelle O, Teo CH (2010) Optimization of ranking measures. *Journal of Machine Learning Research* pp 1–48
- Li P, Burges CJ, Wu Q (2007) McRank: Learning to rank using multiple classification and gradient boosting. In: *Advances in Neural Information Processing Systems (NIPS)*, pp 845–852
- Li Z, Zhang MB, Wang Y, Chen F, Whiffin V, Taib R, Vicky W, Wang Y (2013) Water pipe condition assessment: A hierarchical beta process approach for sparse incident data. *Machine Learning, Special Issue on ML for Science and Society*

- Matveeva I, Laucius A, Burges C, Wong L, Burkard T (2006) High accuracy retrieval with multiple nested ranker. In: Proc. 29th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, pp 437–444
- Menon AK, Jiang X, Kim J, Vaidya J, Ohno-Machado L (2013) Detecting inappropriate access to electronic health records using collaborative filtering. *Machine Learning, Special Issue on ML for Science and Society*
- Metz CE (1978) Basic principles of ROC analysis. *Seminars in Nuclear Medicine* 8(4):283–298
- Oza N, Castle JP, Stutz J (2009) Classification of aeronautics system health and safety documents. *IEEE Transactions on Systems, Man and Cybernetics, Part C* 39:1–11
- Perlich C, Provost F, Simonoff JS (2003) Tree induction vs. logistic regression: A learning-curve analysis. *Journal of Machine Learning Research* 4:211–255
- Potash E, Brew J, Loewi A, Majumdar S, Reece A, Walsh J, Rozier E, Jorgenson E, Mansour R, Ghani R (2015) Predictive modeling for public health: Preventing childhood lead poisoning. In: Proc. 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)
- Putter J (1955) The treatment of ties in some nonparametric tests. *The Annals of Mathematical Statistics* 26(3):368–386
- Qin T, Liu T, Li H (2013) A general approximation framework for direct optimization of information retrieval measures. *Information Retrieval* 13(4):375–397
- Rajaram S, Agarwal S (2005) Generalization bounds for k -partite ranking. In: NIPS 2005 Workshop on Learning to Rank
- Rudin C (2009a) The P-Norm Push: A simple convex ranking algorithm that concentrates at the top of the list. *Journal of Machine Learning Research* 10:2233–2271
- Rudin C (2009b) ROC Flexibility Data. <https://users.cs.duke.edu/~cynthia/code/ROCFlexibilityData.html>
- Rudin C, Schapire RE (2009) Margin-based ranking and an equivalence between AdaBoost and RankBoost. *Journal of Machine Learning Research* 10:2193–2232
- Rudin C, Passonneau R, Radeva A, Dutta H, Jerome S, Isaac D (2010) A process for predicting manhole events in Manhattan. *Machine Learning* 80:1–31
- Rudin C, Waltz D, Anderson RN, Boulanger A, Salleb-Aouissi A, Chow M, Dutta H, Gross P, Huang B, Jerome S, Isaac D, Kressner A, Passonneau RJ, Radeva A, Wu L (2012) Machine learning for the New York City power grid. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34(2):328–345
- Savage IR (1957) Nonparametric statistics. *Journal of the American Statistical Association* 52(279):331–344
- Shen L, Joshi AK (2003) An SVM based voting algorithm with application to parse reranking. In: Proc. HLT-NAACL 2003 workshop on Analysis of Geographic References, pp 9–16
- Tamhane AC, Dunlop DD (2000) *Statistics and data analysis*. Prentice Hall
- Tan M, Xia T, Guo L, Wang S (2013) Direct optimization of ranking measures for learning to rank models. In: Proc. ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), pp 856–864
- Wackerly DD, III WM, Scheaffer RL (2002) *Mathematical statistics with applications*. Duxbury
- Wang Y, Wang L, Li Y, He D, Liu TY (2013) A theoretical analysis of ndcg type ranking measures. In: Proc. 26th Annual Conference on Learning Theory (COLT), PMLR, vol 30, pp 25–54
- Xia F, Liu TY, Wang J, Zhang W, Li H (2008) Listwise approach to learning to rank: theory and algorithm. In: Proc. ICML, ACM, pp 1192–1199
- Xu J (2007) A boosting algorithm for information retrieval. In: Proc. 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval
- Xu J, Li H (2007) Adarank: a boosting algorithm for information retrieval. In: Proc. ACM SIGIR, ACM, pp 391–398
- Yue Y, Finley T, Radlinski F, Joachims T (2007) A support vector method for optimizing average precision. In: Proc. 30th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp 271–278

A Appendix

A.1 Formulation to Maximize Regularized AUC

Again we want to have $z_{ik} = 1$ if $\mathbf{w}^T \mathbf{x}_i > \mathbf{w}^T \mathbf{x}_k$ and $z_{ik} = 0$ otherwise. We want to maximize the sum of the z_{ik} 's which is the number of correctly ranked positive-negative pairs. If $\mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_k \leq \epsilon$ then it is not considered to be correctly ranked. So we need to impose that z_{ik} is 0 when $\mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_k - \epsilon \leq 0$; that is, when 1 plus this quantity is less than 1, z_{ik} is 0. Thus, we impose

$$z_{ik} \leq 1 + \mathbf{w}^T \mathbf{x}_i - \mathbf{w}^T \mathbf{x}_k - \epsilon.$$

Table 5 Detailed experimental results on ROC Flexibility

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	29.12	29.65	31.02	32.34	31.11	30.93	32.27	33.91	32.84	28.95	31.21	1.65
		test	31.52	32.16	33.13	28.64	29.73	31.98	31.59	30.39	30.92	33.45	31.35	1.48
	SVM	train	29.19	29.74	30.83	32.15	31.27	30.84	32.26	32.34	32.81	27.97	30.94	1.57
		test	31.63	32.34	33.09	28.93	29.39	31.94	31.54	28.49	30.92	32.76	31.10	1.63
	RB	train	28.34	27.53	28.04	31.12	30.02	28.63	28.04	30.82	30.14	27.36	29.00	1.39
		test	30.39	30.91	29.98	27.37	28.89	29.69	29.50	27.84	28.90	32.23	29.57	1.43
P-norm Push	train	29.12	29.64	30.85	32.46	31.37	30.93	32.27	33.75	32.81	30.10	31.33	1.48	
	test	31.52	32.14	33.01	29.08	29.39	32.00	31.59	30.25	30.91	34.36	31.43	1.61	
MIO-based methods	$K = 50$	train	31.35	30.60	31.02	33.74	32.78	30.93	32.27	33.91	32.84	30.20	31.96	1.32
		test	33.10	33.44	33.13	30.92	31.70	31.98	31.59	30.39	30.92	34.45	32.16	1.31
	$K = 100$	train	31.44	30.65	30.37	33.82	32.52	30.93	31.71	33.88	32.84	30.28	31.84	1.36
		test	33.12	33.52	32.59	31.01	31.48	31.98	31.29	30.45	30.92	34.51	32.09	1.31
	$K = 150$	train	30.63	30.69	31.06	32.78	32.13	31.27	32.27	32.95	33.00	29.78	31.65	1.12
		test	32.14	33.55	33.19	29.32	30.67	32.33	31.59	29.40	31.10	34.13	31.74	1.65
	Full MIO	train	26.92	27.15	27.04	28.14	30.90	30.99	26.41	30.12	26.96	29.69	28.43	1.80
		test	28.76	29.08	29.16	25.75	28.57	31.95	28.34	27.34	26.74	33.92	28.96	2.40

Regularization is included as usual. The formulation is:

$$\begin{aligned}
& \max_{\mathbf{w}, \gamma_j, z_{ik} \forall j, i, k} \sum_{i \in S_+} \sum_{k \in S_-} z_{ik} - C \sum_j \gamma_j \\
& \text{s.t. } z_{ik} \leq \mathbf{w}^T(\mathbf{x}_i - \mathbf{x}_k) + 1 - \varepsilon, \quad \forall i \in S_+, k \in S_-, \\
& \quad \gamma_j \geq w_j, \quad \forall j = 1, \dots, d, \\
& \quad \gamma_j \geq -w_j, \quad \forall j = 1, \dots, d, \\
& \quad -1 \leq w_j \leq 1, \quad \forall j = 1, \dots, d, \\
& \quad z_{ik}, \gamma_j \in \{0, 1\}, \quad \forall i \in S_+, k \in S_-, j \in \{1, \dots, d\}.
\end{aligned}$$

A.2 Ranking for the General Pairwise Preference Case

RankBoost (Freund et al, 2003) was designed to handle any pairwise preference information. Here we present an exact, regularized version of RankBoost’s objective. Define the labels as $\pi(\mathbf{x}_i, \mathbf{x}_k) = \pi_{ik}$, where π_{ik} is 1 if \mathbf{x}_i should be ranked higher than \mathbf{x}_k . If $\pi_{ik} = 0$ there is no information about the relative ranking of i to k . Then we try to maximize the number of pairs for which the model is able to rank \mathbf{x}_i above \mathbf{x}_k and for which the label for the pair is $\pi_{ik} = 1$:

$$\text{NumAgreedPairs} = \sum_{i=1}^n \sum_{k=1}^n \pi_{ik} \mathbf{1}_{[f(\mathbf{x}_i) > f(\mathbf{x}_k)]}.$$

We will maximize a regularized version of this, as follows:

$$\begin{aligned}
& \max_{\mathbf{w}, \gamma_j, z_{ik} \forall j, i, k} \sum_{i=1}^n \sum_{k=1}^n \pi_{ik} z_{ik} - C \sum_j \gamma_j \\
& \text{s.t. } z_{ik} \leq \mathbf{w}^T(\mathbf{x}_i - \mathbf{x}_k) + 1 - \varepsilon, \quad \forall i, k = 1, \dots, n, \\
& \quad -1 \leq w_j \leq 1, \quad \forall j = 1, \dots, d, \\
& \quad \gamma_j \geq w_j, \quad \forall j = 1, \dots, d, \\
& \quad \gamma_j \geq -w_j, \quad \forall j = 1, \dots, d, \\
& \quad z_{ik}, \gamma_j \in \{0, 1\}, \quad \forall i, k = 1, \dots, n, j \in \{1, \dots, d\}.
\end{aligned}$$

By special choices of π , the pairwise rank statistic can be made to include multipartite ranking (Rajaram and Agarwal, 2005), which can be similar to ordinal regression. In this case we have several classes, where observations in one class should be ranked above (or below) all the observations in another class.

$$\pi_{ik} = \begin{cases} 1 & \text{if observations in Class}(\mathbf{x}_i) \text{ should be ranked above observations in Class}(\mathbf{x}_k), \\ 0 & \text{otherwise.} \end{cases}$$

If there are only two classes, then we are back to the AUC or equivalently the WRS statistic.

A.3 Experimental Results

Table 6 Detailed experimental results on Abalone19

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	3.50	3.56	3.92	3.69	3.56	3.94	2.66	4.19	3.92	3.35	3.63	0.43
		test	3.01	2.91	3.03	3.14	2.75	3.32	3.76	2.22	2.53	2.94	2.96	0.42
	SVM	train	3.46	3.43	3.46	3.39	3.51	2.99	2.46	4.23	3.85	3.37	3.41	0.47
		test	2.94	2.80	3.22	3.18	2.78	3.32	4.22	2.31	2.49	2.92	3.02	0.53
	RB	train	3.73	3.91	2.56	3.32	3.51	2.69	2.34	4.04	3.87	4.06	3.40	0.65
		test	2.47	2.49	2.76	2.89	2.60	3.08	3.66	1.92	2.13	2.62	2.66	0.49
P-norm Push	train	3.44	3.52	3.74	3.40	3.45	3.04	2.44	4.20	3.80	3.33	3.44	0.47	
	test	2.97	2.79	3.05	3.08	2.77	3.23	4.33	2.65	2.55	2.86	3.03	0.50	
MIO-based methods	$K = 250$	train	5.24	4.52	4.82	4.86	5.37	4.58	3.84	5.57	5.69	4.44	4.89	0.58
		test	2.88	3.25	2.99	3.33	3.15	3.22	3.41	2.37	2.26	3.91	3.08	0.49
	$K = 500$	train	4.49	3.72	4.70	4.60	5.10	4.01	3.66	5.04	4.68	4.49	4.45	0.50
		test	2.85	2.87	2.89	3.12	2.50	3.11	3.36	2.20	2.70	3.26	2.89	0.35
	$K = 750$	train	3.79	3.55	4.75	3.94	4.95	3.90	3.57	4.83	3.74	4.27	4.13	0.53
		test	2.64	2.96	2.80	2.90	2.61	3.25	3.14	1.92	2.47	2.96	2.76	0.38
	Full MIO	train	2.60	2.37	2.29	2.54	2.66	2.14	2.00	3.11	2.79	2.89	2.54	0.35
		test	2.22	2.47	2.47	2.98	2.07	2.78	3.24	1.66	1.99	2.32	2.42	0.48

Table 7 Detailed experimental results on UIS

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	18.94	18.64	17.99	19.21	18.50	17.69	21.46	19.38	20.10	16.69	18.86	1.32
		test	17.11	17.84	17.11	16.87	18.72	17.99	17.03	18.15	17.40	20.57	17.88	1.11
	SVM	train	18.37	17.83	17.67	18.39	18.20	17.64	21.17	19.13	20.00	16.18	18.46	1.38
		test	16.91	17.62	17.32	16.78	18.73	17.71	16.91	18.12	17.18	20.80	17.81	1.21
	RB	train	20.40	19.13	19.59	20.40	18.46	18.31	21.97	19.30	20.10	16.69	19.44	1.44
		test	16.18	17.77	18.33	16.30	19.88	17.81	15.64	18.00	17.39	19.67	17.70	1.40
P-norm Push	train	18.93	17.92	17.93	19.18	18.49	17.68	21.55	19.42	20.10	16.60	18.78	1.40	
	test	17.09	17.82	17.19	16.86	18.73	17.96	17.04	18.22	17.42	20.61	17.89	1.13	
MIO-based methods	$K = 50$	train	20.53	20.23	20.52	21.13	19.34	20.32	22.43	20.90	21.25	17.81	20.45	1.23
		test	19.40	19.00	16.57	17.94	18.65	19.73	17.15	16.56	16.14	18.87	18.00	1.31
	$K = 100$	train	20.70	19.66	19.92	20.20	18.33	19.64	21.78	20.04	20.26	17.11	19.76	1.27
		test	18.94	18.69	18.06	18.00	20.60	19.48	15.42	18.93	17.67	20.58	18.64	1.51
	$K = 150$	train	19.91	18.97	19.37	19.53	18.02	19.79	20.32	19.80	20.01	16.90	19.26	1.05
		test	17.59	17.24	18.68	17.11	18.47	17.86	15.64	17.01	16.34	22.00	17.79	1.73
	Full MIO	train	19.06	18.99	18.10	19.90	17.39	18.89	21.68	19.03	19.11	16.24	18.84	1.44
		test	18.47	17.07	18.08	18.03	18.42	18.42	17.02	17.38	17.16	18.82	17.89	0.67

Table 8 Detailed experimental results on Travel

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	31.48	26.63	27.22	28.60	29.25	28.44	27.94	28.70	25.55	27.81	28.16	1.60
		test	23.94	28.69	28.34	26.98	26.76	25.57	27.88	26.95	30.02	28.03	27.32	1.70
	SVM	train	31.08	26.13	26.75	28.23	28.58	27.17	27.29	28.26	25.12	27.29	27.59	1.61
		test	23.41	27.74	27.84	26.93	26.37	24.56	27.22	26.58	29.65	27.79	26.81	1.76
	RB	train	29.83	25.00	25.31	27.01	27.70	27.28	26.72	26.63	24.09	26.20	26.57	1.60
		test	21.93	26.46	26.35	24.33	23.36	24.00	25.20	25.01	27.48	25.40	24.95	1.63
P-norm Push	train	31.52	26.62	27.12	28.49	29.26	28.37	27.93	28.49	25.52	27.61	28.09	1.62	
	test	24.05	28.39	28.25	26.80	26.78	25.51	27.86	26.74	30.00	28.01	27.24	1.66	
MIO-based methods	$K = 50$	train	31.69	26.91	27.18	28.73	29.31	28.74	28.19	28.83	25.58	27.87	28.30	1.63
		test	24.26	29.20	28.48	27.05	26.88	26.42	28.15	26.97	30.46	28.24	27.61	1.70
	$K = 100$	train	31.19	27.05	27.01	28.81	29.62	28.75	27.65	28.53	25.54	28.23	28.24	1.56
		test	24.15	29.63	28.46	27.01	26.31	27.15	26.55	27.48	29.39	27.81	27.39	1.60
	$K = 150$	train	28.76	26.28	25.39	27.95	29.24	27.50	26.90	28.28	24.99	25.93	27.12	1.45
		test	20.99	28.16	26.59	26.53	26.21	23.86	25.99	26.78	29.19	25.68	26.00	2.26
	Full MIO	train	29.57	26.14	26.84	26.55	28.62	26.38	25.85	27.93	25.08	26.44	26.94	1.36
		test	23.16	27.82	27.61	25.01	26.30	24.42	26.41	26.01	29.55	26.85	26.31	1.83

Table 9 Detailed experimental results on NHANES

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	14.50	12.76	16.84	13.89	17.18	15.84	14.62	13.89	15.25	12.14	14.69	1.63
		test	14.43	14.16	10.31	14.63	10.33	11.78	14.74	13.42	12.39	14.40	13.06	1.74
	SVM	train	13.04	11.42	15.80	12.64	16.63	15.72	13.57	13.17	14.99	11.31	13.83	1.87
		test	14.33	14.08	10.27	14.63	10.13	11.75	14.86	13.33	12.12	14.27	12.98	1.79
	RB	train	13.05	12.05	16.73	13.35	16.84	15.20	12.93	13.03	13.87	10.40	13.75	2.01
		test	11.95	13.68	9.78	13.06	9.31	10.67	12.49	12.43	12.67	14.98	12.10	1.75
P-norm Push	train	14.43	12.53	16.63	13.53	16.12	15.89	14.49	13.74	15.32	11.87	14.46	1.57	
	test	14.28	14.39	10.10	14.76	10.24	11.85	14.88	13.63	13.30	14.41	13.18	1.82	
MIO-based methods	$K = 50$	train	15.32	13.04	17.90	14.96	17.84	17.02	14.95	14.64	15.77	13.35	15.48	1.69
		test	14.54	14.39	10.38	13.37	11.07	13.33	14.65	12.99	13.11	14.74	13.26	1.50
	$K = 100$	train	14.11	11.83	17.52	13.77	18.28	16.48	14.49	13.75	15.04	14.95	15.02	1.93
		test	13.58	14.25	10.29	12.76	9.83	12.96	13.53	12.27	14.95	12.74	12.71	1.61
	$K = 150$	train	14.30	12.49	16.91	14.87	16.40	16.56	14.83	13.89	14.77	12.27	14.73	1.59
		test	13.83	14.30	10.73	12.75	10.44	11.91	14.86	13.75	12.40	14.38	12.94	1.55
	Full MIO	train	13.36	12.11	15.79	13.48	15.66	14.48	14.17	13.25	14.13	12.21	13.87	1.25
		test	14.46	14.75	10.42	14.83	10.19	11.76	14.53	13.44	12.19	14.37	13.09	1.82

Table 10 Detailed experimental results on Pima

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	37.27	38.05	33.65	33.84	34.54	33.02	36.10	36.57	35.96	35.98	35.50	1.66
		test	32.64	31.77	36.70	36.07	34.88	36.47	33.85	32.22	34.22	32.94	34.18	1.81
	SVM	train	36.69	37.74	33.52	33.59	34.29	32.93	36.08	36.53	35.88	35.71	35.30	1.61
		test	31.91	31.93	36.73	35.66	34.64	36.41	33.90	32.19	34.24	32.77	34.04	1.83
	RB	train	37.65	36.95	34.29	34.36	35.12	33.27	36.56	36.96	36.29	36.57	35.80	1.44
		test	32.23	31.17	35.97	36.14	34.26	35.40	33.37	33.17	34.02	32.53	33.83	1.65
P-norm Push	train	37.39	38.25	33.74	33.99	34.79	33.11	36.20	36.62	36.10	36.26	35.64	1.67	
	test	32.64	31.70	36.67	36.04	35.10	36.65	33.76	32.21	34.20	33.42	34.24	1.82	
MIO-based methods	$K = 50$	train	37.41	38.36	33.74	34.17	34.56	33.30	36.52	36.71	36.31	36.39	35.75	1.67
		test	32.70	32.18	36.28	36.36	34.91	37.37	33.07	33.86	34.50	33.19	34.44	1.76
	$K = 100$	train	37.58	37.99	33.54	33.67	34.75	33.10	35.10	36.30	36.22	36.07	35.43	1.69
		test	31.41	30.77	33.82	35.93	35.42	35.57	32.34	33.97	33.03	33.31	33.56	1.87
	$K = 150$	train	37.25	37.97	33.00	33.53	34.14	31.35	35.35	36.40	35.54	33.99	34.85	1.96
		test	32.50	31.75	35.93	36.03	35.62	37.05	32.22	32.38	32.98	30.74	33.72	2.21
	Full MIO	train	37.16	37.69	33.78	33.98	34.07	31.99	35.14	36.21	31.71	35.95	34.77	2.03
		test	32.48	31.85	36.06	35.60	35.02	35.19	33.40	32.50	31.10	33.35	33.65	2.01

Table 11 Detailed experimental results on the Gaussians data set

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	69.23	68.26	67.17	68.13	75.38	72.75	68.76	67.18	67.24	68.39	69.25	2.70
		test	64.89	65.62	66.88	65.49	59.25	61.51	64.61	66.58	66.65	65.41	64.89	2.45
	SVM	train	69.30	68.27	67.15	68.19	75.39	72.74	68.90	67.18	67.24	68.38	69.28	2.70
		test	64.92	65.64	66.90	65.59	59.26	61.52	64.81	66.58	66.65	65.41	64.73	2.45
	RB	train	71.14	70.53	69.43	70.51	76.05	74.28	70.21	70.04	69.70	71.22	71.31	2.15
		test	67.58	68.68	68.36	67.81	62.49	64.78	66.91	68.40	69.40	69.36	67.13	2.06
P-norm Push	train	69.22	68.21	67.19	68.12	75.39	72.75	68.69	67.18	67.26	68.38	69.24	2.70	
	test	64.86	65.45	66.86	65.46	59.26	61.51	64.47	66.57	66.65	65.38	64.65	2.43	
MIO-based methods	$K = 50$	train	71.32	70.56	69.36	71.28	76.57	74.59	71.69	70.16	70.03	71.61	71.72	2.22
		test	68.57	69.10	70.38	68.45	62.99	65.26	67.82	69.73	69.81	68.22	68.03	2.27
	$K = 100$	train	71.27	70.43	69.61	71.39	76.56	74.59	71.61	70.33	70.06	71.81	71.76	2.18
		test	68.48	69.12	70.25	68.26	63.05	65.11	67.58	69.81	69.82	67.64	67.91	2.27
	$K = 150$	train	71.01	70.55	69.08	71.41	76.45	74.54	71.76	69.95	69.41	71.67	71.58	2.30
		test	67.31	69.37	70.34	68.46	62.55	65.41	68.01	68.38	69.92	68.11	67.79	2.30
	Full MIO	train	62.81	63.50	61.76	65.16	70.67	66.56	62.53	64.68	64.62	64.69	64.70	2.53
		test	61.28	61.58	62.34	60.50	55.83	55.89	59.76	59.87	61.02	60.81	59.89	2.26

Table 12 Detailed experimental results on Haberman Survival

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	13.45	11.92	14.94	13.98	13.18	13.08	12.47	12.88	12.28	11.17	12.94	1.07
		test	11.45	13.79	10.69	12.01	13.02	13.13	13.47	12.84	13.24	14.58	12.82	1.15
	SVM	train	13.43	11.96	15.02	13.83	13.16	13.19	12.52	12.82	12.34	11.18	12.95	1.06
		test	11.42	13.79	10.70	11.52	13.08	13.06	12.78	12.82	12.58	14.56	12.63	1.15
	RB	train	14.35	13.55	16.48	13.83	13.84	13.81	13.95	14.47	13.89	11.33	13.95	1.24
		test	11.29	13.08	9.60	11.85	12.13	12.62	12.39	11.33	12.74	13.01	12.01	1.06
P-norm Push	train	13.44	11.91	14.88	14.00	13.18	13.11	12.50	12.92	12.26	11.17	12.94	1.06	
	test	11.51	13.66	10.72	11.98	12.94	13.14	12.84	12.80	12.24	14.55	12.64	1.09	
MIO-based	$K = 50$	train	13.68	11.91	14.92	14.12	12.95	13.00	12.54	14.39	12.49	11.02	13.10	1.19
		test	10.10	13.82	10.62	11.99	12.92	13.07	13.48	11.88	13.93	14.60	12.64	1.47
	$K = 100$	train	14.16	12.13	14.94	14.45	12.98	13.03	11.87	12.85	12.63	11.14	13.02	1.20
		test	11.65	13.05	10.52	12.24	12.95	13.19	13.67	12.81	13.23	14.66	12.80	1.13
	Full MIO	train	13.90	12.45	14.92	13.92	12.91	13.26	12.53	12.82	12.97	11.63	13.13	0.92
		test	11.62	13.26	10.60	11.44	12.62	12.82	12.47	12.76	12.61	14.42	12.46	1.05

Table 13 Detailed experimental results on Polypharm

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	17.69	20.69	20.52	19.93	18.63	19.69	21.49	17.99	17.35	20.27	19.43	1.42
		test	18.60	15.65	15.84	16.40	18.60	17.41	14.74	17.82	20.12	17.14	17.23	1.63
	SVM	train	17.57	20.69	20.15	19.85	18.24	19.60	21.46	17.48	17.06	20.26	19.24	1.53
		test	18.51	16.84	17.54	16.93	18.47	17.34	14.67	19.63	20.11	17.09	17.71	1.56
	RB	train	17.65	20.52	20.10	19.97	17.81	19.75	21.03	16.90	17.15	20.25	19.11	1.55
		test	18.98	15.59	16.99	16.10	18.51	17.01	14.89	19.35	19.90	16.67	17.40	1.70
P-norm Push	train	17.30	19.43	19.80	18.71	16.76	19.30	20.53	19.34	18.98	17.16	18.73	1.25	
	test	20.36	17.34	17.75	18.87	19.33	16.96	15.66	17.51	17.95	18.79	18.05	1.33	
MIO-based methods	$K = 50$	train	17.71	19.64	19.95	18.91	19.65	19.48	20.64	19.05	19.37	17.68	19.21	0.93
		test	20.25	16.30	17.02	16.71	18.18	17.58	15.47	17.55	16.49	18.20	17.37	1.33
	$K = 100$	train	17.09	19.29	19.60	18.41	16.89	19.40	20.64	18.07	19.54	17.38	18.63	1.25
		test	19.88	16.82	17.53	18.39	16.96	16.53	15.69	17.40	17.57	19.13	17.59	1.25
	$K = 150$	train	17.11	17.78	17.75	17.77	16.43	19.21	18.99	18.76	19.54	17.05	18.04	1.04
		test	19.93	16.15	13.85	17.70	17.55	17.97	14.21	17.81	16.84	19.54	17.16	1.99
	Full MIO	train	16.95	18.71	19.40	17.21	16.94	18.35	20.33	18.48	19.26	15.92	18.16	1.36
		test	20.05	16.50	17.41	17.68	18.73	14.62	15.84	16.92	16.16	18.28	17.22	1.57

Table 14 Detailed experimental results on Glow500

Algorithm		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
Baseline methods	LR	train	14.66	17.70	18.65	17.44	18.44	17.84	16.95	16.51	16.57	17.87	17.26	1.16
		test	18.32	15.45	16.06	17.11	14.78	17.68	17.64	16.90	17.45	15.43	16.68	1.18
	SVM	train	13.94	17.13	18.67	16.92	17.57	17.80	16.18	15.41	15.07	16.52	16.52	1.41
		test	18.48	15.83	15.59	17.11	16.31	17.63	18.21	17.90	18.91	16.59	17.26	1.15
	RB	train	14.34	17.77	18.65	18.12	17.91	17.22	17.94	16.31	16.42	17.65	17.23	1.25
		test	18.48	16.44	16.77	16.29	16.32	17.71	16.83	17.37	17.20	16.67	17.01	0.69
P-norm Push	train	14.10	17.67	18.96	17.36	18.52	17.97	16.56	15.66	16.19	17.22	17.02	1.44	
	test	18.35	15.92	15.68	17.56	15.41	18.07	18.65	18.01	18.45	16.28	17.24	1.27	
MIO-based methods	$K = 50$	train	15.62	18.85	19.20	19.02	19.21	18.65	18.30	17.31	16.74	19.12	18.20	1.24
		test	16.89	15.23	17.18	16.43	17.03	18.33	17.65	17.30	17.15	14.68	16.79	1.09
	$K = 100$	train	15.30	18.27	18.51	18.32	19.42	18.30	17.86	16.18	16.93	17.89	17.70	1.22
		test	17.38	16.82	15.39	17.69	16.18	18.50	19.06	15.74	19.68	14.59	17.10	1.66
	$K = 150$	train	14.53	17.85	18.69	17.90	18.81	17.25	16.87	15.85	15.53	17.53	17.08	1.40
		test	19.26	14.85	15.58	16.73	16.07	17.40	17.27	18.28	16.70	15.65	16.78	1.33
	Full MIO	train	13.51	17.51	17.46	16.46	18.40	16.97	16.71	16.35	16.64	16.94	16.69	1.28
		test	16.47	15.59	15.74	15.78	15.54	17.26	15.72	18.12	19.36	15.22	16.48	1.35

Table 15 Detailed experimental results on ROC Flexibility, with different c and ϵ values.

Parameter values		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
$K = 100, \epsilon = 10^{-4}$	$c = 10^{-1}$	train	29.12	29.65	31.02	32.34	31.11	30.93	32.27	33.91	32.84	29.95	31.31	1.52
		test	31.52	32.16	33.13	28.64	29.73	31.98	31.59	30.39	30.92	33.45	31.35	1.48
	$c = 10^{-2}$	train	30.35	30.18	27.91	33.80	31.82	30.93	32.27	32.71	32.84	30.28	31.31	1.72
		test	31.78	32.42	29.64	30.97	30.35	31.98	31.59	28.83	30.92	34.51	31.30	1.57
	$c = 10^{-3}$	train	31.44	30.65	30.37	33.82	32.52	30.93	31.71	33.88	32.84	30.28	31.84	1.36
		test	33.12	33.52	32.59	31.01	31.48	31.98	31.29	30.45	30.92	34.51	32.09	1.31
	$c = 10^{-4}$	train	31.50	30.69	31.06	33.14	32.74	30.93	32.27	33.93	32.84	30.35	31.94	1.20
		test	33.17	33.55	33.19	29.59	31.65	31.98	31.59	30.48	30.92	34.52	32.06	1.53
	$c = 10^{-5}$	train	31.50	30.69	31.06	33.86	32.74	30.93	32.27	33.93	32.84	30.35	32.02	1.30
		test	33.17	33.55	33.19	31.08	31.65	31.98	31.59	30.48	30.92	34.52	32.21	1.32
	$c = 10^{-6}$	train	30.63	30.60	31.06	33.86	32.13	30.93	32.27	33.93	32.84	30.35	31.86	1.35
		test	32.14	33.55	33.19	31.08	30.67	31.98	31.59	30.48	30.92	34.52	32.01	1.35
$K = 100, c = 10^{-3}$	$\epsilon = 10^{-1}$	train	31.45	30.69	31.06	33.05	30.08	30.93	32.27	33.65	32.84	30.24	31.62	1.25
		test	32.12	33.55	33.19	29.21	28.21	31.98	31.59	29.62	30.92	34.47	31.59	2.07
	$\epsilon = 10^{-2}$	train	30.63	30.69	31.06	33.05	32.74	30.93	32.27	33.92	32.84	30.35	31.85	1.26
		test	32.14	33.55	33.19	29.11	31.65	31.98	31.59	30.49	30.92	34.52	31.91	1.57
	$\epsilon = 10^{-3}$	train	30.63	30.69	31.06	33.84	32.74	30.93	32.27	33.93	32.84	30.35	31.93	1.36
		test	32.14	33.55	33.19	31.03	31.65	31.98	31.59	30.48	30.92	34.52	32.10	1.28
	$\epsilon = 10^{-4}$	train	31.44	30.65	30.37	33.82	32.52	30.93	31.71	33.88	32.84	30.28	31.84	1.36
		test	33.12	33.52	32.59	31.01	31.48	31.98	31.29	30.45	30.92	34.51	32.09	1.31
	$\epsilon = 10^{-5}$	train	31.50	30.69	31.06	33.86	32.74	30.93	32.27	33.93	32.84	30.35	32.02	1.30
		test	33.17	33.55	33.19	31.08	31.65	31.98	31.59	30.48	30.92	34.52	32.21	1.32
	$\epsilon = 10^{-6}$	train	31.38	30.19	31.06	33.74	31.55	30.93	31.93	33.90	32.84	28.27	31.58	1.68
		test	33.07	33.04	33.19	30.92	29.61	31.98	31.48	30.41	30.92	32.75	31.74	1.26

Table 16 Detailed experimental results on UIS, with different c and ε values.

Parameter values		Runs										Statistics		
		1	2	3	4	5	6	7	8	9	10	Mean	Std. Dev.	
$K = 100, \varepsilon = 10^{-4}$	$c = 10^{-1}$	train	18.94	18.64	17.99	19.21	18.50	20.25	21.46	19.38	20.10	17.08	19.15	1.24
		test	17.11	17.84	17.11	16.87	18.72	18.72	17.03	18.15	17.40	20.44	17.94	1.11
	$c = 10^{-2}$	train	19.39	18.86	19.75	19.85	18.56	19.45	21.40	19.32	19.63	17.39	19.36	1.02
		test	18.21	18.39	17.14	17.84	20.71	16.78	16.05	18.44	17.20	20.70	18.15	1.54
	$c = 10^{-3}$	train	20.17	19.92	19.47	19.80	18.81	20.16	21.46	19.82	19.71	17.53	19.69	1.01
		test	18.44	18.30	17.43	17.23	19.83	18.73	16.15	16.09	16.31	20.67	17.92	1.57
	$c = 10^{-4}$	train	20.70	19.66	19.92	20.20	18.33	19.64	21.78	20.04	20.26	17.11	19.76	1.27
		test	18.94	18.69	18.06	18.00	20.60	19.48	15.42	18.93	17.67	20.58	18.64	1.51
	$c = 10^{-5}$	train	20.55	19.03	19.73	20.83	19.00	19.46	22.16	20.02	20.48	17.61	19.89	1.24
		test	18.84	18.71	16.98	17.49	18.53	18.17	16.67	17.21	16.73	20.12	17.94	1.12
	$c = 10^{-6}$	train	20.67	19.95	19.51	19.83	18.88	19.22	21.70	19.76	19.70	17.55	19.68	1.08
		test	18.41	17.52	16.80	17.90	18.30	19.93	15.38	15.92	16.25	20.96	17.74	1.76
$K = 100, c = 10^{-4}$	$\varepsilon = 10^{-1}$	train	20.22	19.26	19.37	20.45	18.86	19.90	21.87	19.88	20.08	17.10	19.70	1.23
		test	18.78	17.65	17.07	17.80	18.78	19.71	17.22	18.38	18.11	20.53	18.40	1.09
	$\varepsilon = 10^{-2}$	train	20.22	19.61	19.64	20.43	18.15	19.54	22.17	19.88	20.04	17.61	19.73	1.24
		test	18.28	18.56	18.45	17.39	19.47	17.11	16.71	18.42	16.46	19.44	18.03	1.06
	$\varepsilon = 10^{-3}$	train	20.97	19.30	20.06	19.65	18.61	20.07	21.35	20.10	20.26	17.63	19.80	1.09
		test	17.98	18.00	18.01	17.46	20.04	18.67	16.45	17.91	18.26	20.60	18.34	1.20
	$\varepsilon = 10^{-4}$	train	20.70	19.66	19.92	20.20	18.33	19.64	21.78	20.04	20.26	17.11	19.76	1.27
		test	18.94	18.69	18.06	18.00	20.60	19.48	15.42	18.93	17.67	20.58	18.64	1.51
	$\varepsilon = 10^{-5}$	train	20.43	19.81	19.53	19.98	18.47	20.09	21.55	19.95	20.32	17.20	19.73	1.17
		test	17.06	18.74	16.70	17.95	19.08	19.08	15.03	17.83	16.45	20.86	17.88	1.66
	$\varepsilon = 10^{-6}$	train	20.16	18.37	19.53	18.15	18.50	18.81	21.46	19.38	20.10	16.32	19.08	1.40
		test	18.23	18.24	17.66	18.23	18.72	19.51	17.03	18.15	17.40	19.12	18.23	0.75