

# Active Sampling of Pairs and Points for Large-scale Linear Bipartite Ranking

Wei-Yuan Shen

Hsuan-Tien Lin

*Department of Computer Science and Information Engineering  
National Taiwan University*

R00922024@CSIE.NTU.EDU.TW

HTLIN@CSIE.NTU.EDU.TW

## Abstract

Bipartite ranking is a fundamental ranking problem that learns to order relevant instances ahead of irrelevant ones. One major approach for bipartite ranking, called the pair-wise approach, tackles an equivalent binary classification problem of whether one instance out of a pair of instances should be ranked higher than the other. Nevertheless, the number of instance pairs constructed from the input data could be quadratic to the size of the input data, which makes pair-wise ranking generally infeasible on large-scale data sets. Another major approach for bipartite ranking, called the point-wise approach, directly solves a binary classification problem between relevant and irrelevant instance points. This approach is feasible for large-scale data sets, but the resulting ranking performance can be inferior. That is, it is difficult to conduct bipartite ranking accurately and efficiently at the same time. In this paper, we develop a novel scheme within the pair-wise approach to conduct bipartite ranking efficiently. The scheme, called Active Sampling, is inspired from the rich field of active learning and can reach a competitive ranking performance while focusing only on a small subset of the many pairs during training. Moreover, we propose a general Combined Ranking and Classification (CRC) framework to accurately conduct bipartite ranking. The framework unifies point-wise and pair-wise approaches and is simply based on the idea of treating each instance point as a pseudo-pair. Experiments on 14 real-world large-scale data sets demonstrate that the proposed algorithm of Active Sampling within CRC, when coupled with a linear Support Vector Machine, usually outperforms state-of-the-art point-wise and pair-wise ranking approaches in terms of both accuracy and efficiency.

**Keywords:** bipartite ranking, binary classification, large-scale, active learning, AUC.

## 1. Introduction

The bipartite ranking problem aims at learning a ranking function that orders positive instances ahead of negative ones. For example, in information retrieval, bipartite ranking can be used to order the preferred documents in front of the less-preferred ones within a list of search-engine results. The performance of the ranking function is measured by the probability of mis-ordering an unseen pair of randomly chosen positive and negative instances, which is equal to one minus the Area Under the ROC Curve (AUC) [14], a popular criterion for evaluating the sensitivity and the specificity of binary classifiers in many real-world tasks [5] and large-scale data mining competitions [7].

Given the many potential applications in information retrieval, bioinformatics, and recommendation systems, bipartite ranking has received much research attention in the past two decades [1; 5; 9; 12; 15; 20; 22; 24]. Many existing bipartite ranking algorithms explicitly or implicitly reduce the problem to binary classification to inherit the benefits from the well-developed methods in binary classification [5; 12; 15; 18; 22]. The majority of those reduction-based algorithms can be categorized to two approaches: the pair-wise approach and the point-wise one. The pair-wise approach transforms the input data of positive and negative instances to *pairs* of instances, and learns a binary classifier for predicting whether the first instance in a pair should be scored higher than the second one. The pair-wise approach comes with strong theoretical guarantee. For example, [3] shows that a low-regret ranking function can indeed be formed by a low-regret binary classifier. The strong theoretical guarantee leads to promising experimental results in many state-of-the-art bipartite ranking algorithms, such as RankSVM [18], RankBoost [15] and RankNet [6]. Nevertheless, the number of pairs in the input data can easily be of size  $\Theta(N^2)$ , where  $N$  is the size of the input data, if the data is not extremely unbalanced. The quadratic number of pairs with respect to  $N$  makes the pair-wise approach computationally infeasible for large-scale data sets in general, except in a few special algorithms like RankBoost [15] or the efficient linear RankSVM [20]. RankBoost enjoys an efficient implementation by reducing the quadratic number of pair-wise terms in the objective function to a linear number of equivalent terms; efficient linear RankSVM transforms the pair-wise optimization formulation to an equivalent formulation that can be solved in subquadratic time complexity [22].

On the other hand, the point-wise approach directly runs binary classification on the positive and negative instance *points* of the input data, and takes the scoring function behind the resulting binary classifier as the ranking function. In some special cases [15], such as AdaBoost [16] and its pair-wise sibling RankBoost [15], the point-wise approach is shown to be equivalent to the corresponding pair-wise one [12]. In other cases, the point-wise approach often operates with an approximate objective function that involves only  $N$  terms [22]. For example, [22] shows that minimizing the exponential or the logistic loss function on the instance points decreases an upper bound on the number of mis-ordered pairs within the input data. Because of the approximate nature of the point-wise approach, its ranking performance can sometimes be inferior to the pair-wise approach.

From the discussion above, we see that the pair-wise approach leads to more satisfactory performance while the point-wise approach comes with efficiency, and there is a trade-off between the two. In this paper, we are interested in designing bipartite ranking algorithms that enjoy both satisfactory performance and efficiency for large-scale bipartite ranking. We focus on using the linear Support Vector Machine (SVM) [31] given its recent advances for efficient large-scale learning [33]. We first show that the loss function behind the usual point-wise SVM [31] minimizes an upper bound on the loss function behind RankSVM, which suggests that the point-wise SVM could be an approximate bipartite ranking algorithm that enjoys efficiency. Then, we design a better ranking algorithm with two major contributions.

Firstly, we study an active sampling scheme to select important pairs for the pair-wise approach and name the scheme Active Sampling for RankSVM (ASRankSVM). The scheme makes the pair-wise SVM computationally feasible by focusing only on a small number of valuable pairs out of the quadratic number of pairs. The active sampling scheme is inspired by active learning, another popular machine learning setup that aims to save the efforts

of labeling [27]. More specifically, we discuss the similarity and differences between active sampling and pool-based active learning, and propose some active sampling strategies based on the similarity. Secondly, we propose a general framework that unifies the point-wise SVM and the pair-wise SVM (RankSVM) as special cases. The framework, called combined ranking and classification (CRC), is simply based on the idea of treating each instance point as a pseudo-pair. The CRC framework coupled with active sampling improves the performance of the point-wise SVM by considering not only points but also pairs in its objective function.

Performing active sampling within the CRC framework leads to a promising algorithm for large-scale linear bipartite ranking. We conduct experiments on 14 real-world large-scale data sets and compare the proposed algorithms (ASRankSVM and ASCRC) with several state-of-the-art bipartite ranking algorithms, including the point-wise linear SVM [13], the efficient linear RankSVM [20], and the Combined Ranking and Regression (CRR) algorithm [26] which is closely related to the CRC framework. The results show that ASRankSVM is able to efficiently sample only 8,000 of the more than millions of the possible pairs to achieve better performance than other state-of-the-art algorithms that use all the pairs, while ASCRC that considers the pseudo-pairs can sometimes be helpful. Those results validate that the proposed algorithm can indeed enjoy both satisfactory performance and efficiency for large-scale bipartite ranking.

The paper is organized as follows. Section 2 describe the problem setup and several related works in the literature. Then, we illustrate the active sampling scheme and the CRC framework in Section 3. We conduct a thorough experimental study to compare the proposed algorithm to several state-of-the-art ones in Section 4, and conclude in Section 5.

## 2. Setup and Related Works

In a bipartite ranking problem, we are given a training set  $\mathcal{D} = \{(\mathbf{x}_k, y_k)\}_{k=1}^N$ , where each  $(\mathbf{x}_k, y_k)$  is a training instance with the feature vector  $\mathbf{x}_k$  in an  $n$ -dimensional space  $\mathcal{X} \subseteq \mathbb{R}^n$  and the binary label  $y_k \in \{+1, -1\}$ . Such a training set is of the same format as the training set in usual binary classification problems. We assume that the instances  $(\mathbf{x}_k, y_k)$  are drawn *i.i.d.* from an unknown distribution  $P$  on  $\mathcal{X} \times \{+1, -1\}$ . Bipartite ranking algorithms take  $\mathcal{D}$  as the input and learn a ranking function  $r: \mathcal{X} \rightarrow \mathbb{R}$  that maps a feature vector  $\mathbf{x}$  to a real-valued score  $r(\mathbf{x})$ .

For any pair of two instances, we call the pair mis-ordered by  $r$  *iff* the pair contains a positive instance  $(\mathbf{x}_+, +1)$  and a negative one  $(\mathbf{x}_-, -1)$  while  $r(\mathbf{x}_+) \leq r(\mathbf{x}_-)$ . For a distribution  $P$  that generates instances  $(\mathbf{x}, y)$ , we can define its pair distribution  $P_2$  that generates  $(\mathbf{x}, y, \mathbf{x}', y')$  to be the conditional probability of sampling two instances  $(\mathbf{x}, y)$  and  $(\mathbf{x}', y')$  from  $P$ , conditioned on  $y \neq y'$ . Then, let the expected bipartite ranking loss  $L_P(r)$  for any ranking function  $r$  be the expected number of mis-ordered pairs over  $P_2$ .

$$L_P(r) = \mathbb{E}_{(\mathbf{x}, y, \mathbf{x}', y') \sim P_2} \left[ \mathbb{I} \left( (y - y')(r(\mathbf{x}) - r(\mathbf{x}')) \leq 0 \right) \right],$$

where  $\mathbb{I}(\bullet)$  is an indicator function that returns 1 *iff* the condition  $(\bullet)$  is true, and returns 0 otherwise. The goal of bipartite ranking is to use the training set  $\mathcal{D}$  to learn a ranking function  $r$  that minimizes the expected bipartite ranking loss  $L_P(r)$ . Because  $P$  is unknown,

$L_P(r)$  cannot be computed directly. Thus, bipartite ranking algorithms usually resort to the empirical bipartite ranking loss  $L_{\mathcal{D}}(r)$ , which takes the expectation over the pairs in  $\mathcal{D}$  instead of over the pair distribution  $P_2$ .

The bipartite ranking loss  $L_P(r)$  is closely related to the area under the ROC curve (AUC), which calculates the expected number of *correctly-ordered* pairs. Hence,  $\text{AUC}_{\bullet}(r) = 1 - L_{\bullet}(r)$  for  $\bullet = P$  or  $\mathcal{D}$ , and higher AUC indicates better ranking performance. Bipartite ranking is a special case of the general ranking problem in which the labels  $y$  can be any real value, not necessarily  $\{+1, -1\}$ .

Motivated by the recent advances of linear models for efficient large-scale learning [33], we consider linear models for efficient large-scale bipartite ranking. That is, the ranking functions would be of the form  $r(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$ . In particular, we study the linear Support Vector Machine (SVM) [31] for bipartite ranking. There are two possible approaches for adopting the linear SVM on bipartite ranking problems, the *pair-wise* SVM approach and the *point-wise* SVM approach.

The pair-wise approach corresponds to the famous RankSVM algorithm [18], which is originally designed for ranking with ordinal-scaled scores, but can be easily extended to general ranking with real-valued labels or restricted to bipartite ranking with binary labels. For each positive instance  $(\mathbf{x}_i, y_i = +1)$  and negative instance  $(\mathbf{x}_j, y_j = -1)$ , the pair-wise approach transforms the two instances to two symmetric *pairs* of instance  $((\mathbf{x}_i, \mathbf{x}_j), +1)$  and  $((\mathbf{x}_j, \mathbf{x}_i), -1)$ , the former for indicating that  $\mathbf{x}_i$  should be scored higher than  $\mathbf{x}_j$  and the latter for indicating that  $\mathbf{x}_j$  should be scored lower than  $\mathbf{x}_i$ . The pairs transformed from  $\mathcal{D}$  are then fed to an SVM for learning a ranking function of the form  $r(\mathbf{x}) = \mathbf{w}^T \phi(\mathbf{x})$ , where  $\phi$  indicates some feature transform.

When using a linear SVM,  $\phi$  is simply the identity function. Then, for the pair  $((\mathbf{x}_i, \mathbf{x}_j), +1)$ , we see that  $\mathbb{I}(r(\mathbf{x}_i) \leq r(\mathbf{x}_j)) = 0$  iff  $\mathbf{w}^T(\mathbf{x}_i - \mathbf{x}_j) > 0$ . Define the transformed feature vector  $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$  and the transformed label  $y_{ij} = \text{sign}(y_i - y_j)$ , we can equivalently view the pair-wise linear SVM as simply running a linear SVM on the pair-wise training set  $\mathcal{D}_{pair} = \{(\mathbf{x}_{ij}, y_{ij}) | y_i \neq y_j\}$ . The pair-wise linear SVM minimizes the hinge loss as a surrogate to the 0/1 loss on  $\mathcal{D}_{pair}$  [29], and the 0/1 loss on  $\mathcal{D}_{pair}$  is equivalent to  $L_{\mathcal{D}}(r)$ , the empirical bipartite ranking loss of interest. That is, if the linear SVM learns an accurate binary classifier using  $\mathcal{D}_{pair}$ , the resulting ranker  $r(\mathbf{x}) = \mathbf{w}^T \mathbf{x}$  would also be accurate in terms of the bipartite ranking loss.

Denote the hinge function  $\max(\bullet, 0)$  by  $[\bullet]_+$ , RankSVM solves the following optimization problem

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{\mathbf{x}_{ij} \in \mathcal{D}_{pair}} C_{ij} [1 - \mathbf{w}^T y_{ij} \mathbf{x}_{ij}]_+, \quad (1)$$

where  $C_{ij}$  denotes the weight of the pair  $\mathbf{x}_{ij}$ . Because of the symmetry of  $\mathbf{x}_{ij}$  and  $\mathbf{x}_{ji}$ , we naturally assume that  $C_{ij} = C_{ji}$ . In the original RankSVM formulation,  $C_{ij}$  is set to a constant for all the pairs. Here we list a more flexible formulation (1) to facilitate some discussions later. RankSVM has reached promising bipartite ranking performance in the literature [5]. Because of the symmetry of positive and negative pairs, we can equivalently solve (1) on those positive pairs with  $y_{ij} = 1$ . The number of such positive pairs is  $N^+ N^-$  if there are  $N^+$  positive instances and  $N^-$  negative ones. The huge number of pairs make it difficult to solve (1) with a naïve quadratic programming algorithm.

In contrast with the naïve RankSVM, the efficient linear RankSVM [20] changes (1) to a more sophisticated and equivalent one with an exponential number of constraints, each corresponding to a particular linear combination of the pairs. Then, it reaches  $O(N \log N)$  time complexity by using a cutting-plane solver to identify the most-violated constraints iteratively, while the constant hidden in the big- $O$  notation depends on the parameter  $C_{ij}$  as well as the desired precision of optimization. The subquadratic time complexity of the efficient RankSVM can make it much slower than the point-wise approach (to be discussed below), and hence may not always be fast enough for large-scale bipartite ranking.

The *point-wise* SVM approach, on the other hand, directly runs an SVM on the original training set  $\mathcal{D}$  instead of  $\mathcal{D}_{pair}$ . That is, in the linear case, the point-wise approach solves the following optimization problem

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{w} + C_+ \sum_{\mathbf{x}_i \in \mathcal{D}^+} [1 - \mathbf{w}^T \mathbf{x}_i]_+ + C_- \sum_{\mathbf{x}_j \in \mathcal{D}^-} [1 + \mathbf{w}^T \mathbf{x}_j]_+ . \quad (2)$$

Such an approach comes with some theoretical justification [22]. In particular, the 0/1 loss on  $\mathcal{D}$  has been proved to be an upper bound of the empirical bipartite ranking loss. In fact, the bound can be tightened by adjusting  $C_+$  and  $C_-$  to balance the distribution of the positive and negative instances in  $\mathcal{D}$ . When  $C_+ = C_-$ , [5] shows that the point-wise approach (2) is inferior to the pair-wise approach (1) in performance. The inferior performance can be attributed to the fact that the point-wise approach only operates with an approximation (upper bound) of the bipartite ranking loss of interest.

Next, inspired by the theoretical result of upper-bounding the bipartite ranking loss with a balanced 0/1 loss, we study the connection between (1) and (2) by balancing the hinge loss in (2). In particular, as shown in Theorem 1, a balanced form of (2) can be viewed as minimizing an upper bound of the objective function within (1). In other words, the weighted point-wise SVM can be viewed as a reasonable baseline algorithm for large-scale bipartite ranking problem.

**Theorem 1** *Let  $C_{ij} = \frac{C}{2}$  be a constant in (1);  $C_+ = 2N^- \cdot C$  and  $C_- = 2N^+ \cdot C$  in (2). Then, the objective function of (1) is upper-bounded by  $\frac{1}{4}$  times the objective function of (2).*

**Proof** Because  $[1 - \mathbf{w}^T \mathbf{x}_{ij}]_+ \leq \frac{1}{2} ([1 - 2\mathbf{w}^T \mathbf{x}_i]_+ + [1 + 2\mathbf{w}^T \mathbf{x}_j]_+)$ ,

$$\begin{aligned} & \frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{\mathbf{x}_{ij} \in \mathcal{D}_{pair}} C_{ij} [1 - \mathbf{w}^T y_{ij} \mathbf{x}_{ij}]_+ = \frac{1}{2} \mathbf{w}^T \mathbf{w} + \sum_{\mathbf{x}_{ij} \in \mathcal{D}_{pair}, y_{ij}=+1} C [1 - \mathbf{w}^T \mathbf{x}_{ij}]_+ \\ & \leq \frac{1}{2} \mathbf{w}^T \mathbf{w} + \frac{C}{2} \sum_{\mathbf{x}_i \in \mathcal{D}^+} \sum_{\mathbf{x}_j \in \mathcal{D}^-} ([1 - 2\mathbf{w}^T \mathbf{x}_i]_+ + [1 + 2\mathbf{w}^T \mathbf{x}_j]_+) \\ & = \frac{1}{2} \mathbf{w}^T \mathbf{w} + \frac{C}{2} \left( N^- \sum_{\mathbf{x}_i \in \mathcal{D}^+} [1 - 2\mathbf{w}^T \mathbf{x}_i]_+ + N^+ \sum_{\mathbf{x}_j \in \mathcal{D}^-} [1 + 2\mathbf{w}^T \mathbf{x}_j]_+ \right) . \end{aligned}$$

The theorem can be easily proved by substituting  $2\mathbf{w}$  with a new variable  $\mathbf{u}$ . ■

### 3. Bipartite Ranking with Active Querying

As discussed in the previous section, the pair-wise approach (1) is infeasible on large-scale data sets due to the huge number of pairs. Then, either some random sub-sampling of the pairs are needed [26], or the less-accurate point-wise approach (2) is taken as the approximate alternative [22]. Nevertheless, the better ranking performance of the pair-wise approach over the point-wise one suggest that some of the key pairs shall carry more valuable information than the instance-points. Next, we design an algorithm that samples a few key pairs actively during learning. The resulting algorithm achieves better performance than the point-wise approaches because of the key pairs, and enjoys better efficiency than the pair-wise approach because of the sampling. We first show that some proposed active sampling schemes, which are inspired by the many existing methods in active learning [23; 25; 27], can help identify those key pairs better than random sub-sampling. Then, we discuss how we can unify point-wise and pair-wise ranking approaches under the same framework.

#### 3.1. Pool-based Active Learning

The pair-wise SVM approach (1) is challenging to solve because of the huge number of pairs involved in  $\mathcal{D}_{pair}$ . To make the computation feasible, we can only afford to work on a small subset of  $\mathcal{D}_{pair}$  during training. Existing algorithms conquer the computational difficulty of the huge number of pairs in different ways. The Combined Ranking and Regression approach [26] performs stochastic gradient descent, which essentially selects within the huge number of pairs in a random manner; the efficient RankSVM [20] identifies the most-violated constraints during optimization, which corresponds to selecting the most valuable pairs from an optimization perspective.

We take an alternative route and hope to select the most valuable pairs from a *learning* perspective. That is, our task is to iteratively select a small number of valuable pairs for training while reaching similar performance to the pair-wise approach that trains with all the pairs. One machine learning setup that works for a similar task is *active learning* [27], which iteratively select a small number of valuable instances for labeling (and training) while reaching similar performance to the approach that trains with all the instances fully labeled. [2] proves that selecting a subquadratic number of pairs is sufficient to obtain a ranking function that is close to the optimal ranking function. The algorithm is theoretical in nature, while many other promising active learning tools [23; 25; 27] have not been explored for selecting valuable pairs in large-scale bipartite ranking.

Next, we start exploring those tools by providing a brief review about active learning. We focus on the setup of pool-based active learning [27] because of its strong connection to our needs. In a pool-based active learning problem, the training instances are separated into two parts, the labeled pool ( $\mathcal{L}$ ) and the unlabeled pool ( $\mathcal{U}$ ). As the name suggests, the labeled pool consists of labeled instances that contain both the feature vector  $\mathbf{x}_k$  and its corresponding label  $y_k$ , and the unlabeled pool contains unlabeled instances  $\mathbf{x}_\ell$  only. Pool-based active learning assumes that a (huge) pool of unlabeled instances is relatively easy to gather, while labeling those instances can be expensive. Therefore, we hope to achieve promising learning performance with as few labeled instances as possible. A pool-based active learning algorithm is generally iterative. In each iteration, there are two steps: the training step and the querying step. In the training step, the algorithm trains a decision

function from the labeled pool; in the querying step, the algorithm selects one (or a few) unlabeled instances, queries an oracle to label those instances, and move those instances from the unlabeled pool to the labeled one. The pool-based active learning framework repeats the training and querying steps iteratively until a given budget  $B$  on the number of queries is met, with the hope that the decision functions returned throughout the learning steps are as accurate as possible for prediction.

Because labeling is expensive, active learning algorithms aim to select the most valuable instance(s) from the unlabeled pool at each querying step. Various selection criteria have been proposed to describe the value of an unlabeled instance [27], such as uncertainty sampling [23], and expected error reduction [25].

Moreover, there are several works that solve bipartite ranking under the active learning scenario [10; 11; 32]. For example, [10] selects points that reduce the ranking loss functions most from the unlabeled pool while [11] selects points that maximize the AUC in expectation. Nevertheless, these active learning algorithms require either sorting or enumerating over the huge unlabeled pool in each querying step. The sorting or enumerating process can be time consuming, but have not been considered seriously because labeling is assumed to be even more expensive. We will discuss later that those algorithms that require sorting or enumerating may not fit our goal.

### 3.2. Active Sampling

Following the philosophy of active learning, we propose the Active Sampling scheme for choosing a smaller set of key pairs on the huge training set  $\mathcal{D}_{pair}$ . We call the scheme *Active Sampling* in order to highlight some differences to active learning. One particular difference is that RankSVM (1) only requires optimizing with *positive pairs*. Then, the label  $y_{ij}$  of a pair is a constant 1 and thus easy to get during Active Sampling, while the label in active learning remains unknown before the possibly expensive querying step. Thus, while Active Sampling and active learning both focus on using as few labeled data as possible, the costly part of the Active Sampling scheme is on training rather than querying.

For Active Sampling, we denote  $B$  as the budget on the number of pairs that can be used in training, which plays a similar role to the budget on querying in active learning. We separate the pair-wise training set  $\mathcal{D}_{pair}$  into two parts, the chosen pool ( $\mathcal{L}^*$ ) and the unchosen pool ( $\mathcal{U}^*$ ). The chosen pool is the subset of pairs to be used for training, and the unchosen pool contains the unused pairs. The chosen pool is similar to the labeled pool in pool-based active learning; the unchosen pool acts like the unlabeled pool. The fact that it is almost costless to “label” the instances in the unchosen pool allows us to design simpler sampling strategies than those commonly used for active learning, because no effort is needed to estimate the unknown labels.

The proposed scheme of Active Sampling is illustrated in Algorithm 1. The algorithm takes an initial chosen pool  $\mathcal{L}^*$  and an initial unchosen pool  $\mathcal{U}^*$ , where we simply mimic the usual setup in pool-based active learning by letting  $\mathcal{L}^*$  be a randomly chosen subset of  $\mathcal{D}_{pair}$  and  $\mathcal{U}^*$  be the set of unchosen pairs in  $\mathcal{D}_{pair}$ . In each iteration of the algorithm, we use `Sample` to actively choose  $b$  instances to be moved from  $\mathcal{U}^*$  to  $\mathcal{L}^*$ . After sampling, a `linearSVM` is called to learn from  $\mathcal{L}^*$  along with the weights in  $\{C_{ij}\}$ . We feed the current  $\mathbf{w}$  to the `linearSVM` solver to allow a *warm-start* in optimization. The warm-start step

**Input:** the initial chosen pool,  $\mathcal{L}^*$ ; the initial unchosen pool,  $\mathcal{U}^*$ ; the regularization parameters,  $\{C_{ij}\}$ . the number of pairs sampled per iteration,  $b$ ; the budget on the total number of pairs sampled,  $B$ ; the sampling strategy,  $\text{Sample}(\mathcal{U}^*, \mathbf{w}) \rightarrow \mathbf{x}_{ij}$  that chooses a pair from  $\mathcal{U}^*$ .

**Output:** the ranking function represented by the weights  $\mathbf{w}$ .

$\mathbf{w} = \text{linearSVM}(\mathcal{L}^*, \{C_{ij}\}, \mathbf{0})$  ;

**repeat**

**for**  $i = 1 \rightarrow b$  **do**

$\mathbf{x}_{ij} = \text{Sample}(\mathcal{U}^*, \mathbf{w})$ ;

$\mathcal{L}^* = \mathcal{L}^* \cup \{(\mathbf{x}_{ij}, y_{ij})\}$  and  $\mathcal{U}^* = \mathcal{U}^* \setminus \{\mathbf{x}_{ij}\}$ ;

**end**

$\mathbf{w} = \text{linearSVM}(\mathcal{L}^*, \{C_{ij}\}, \mathbf{w})$ ;

**until**  $(|\mathcal{L}^*| \geq B)$ ;

**return**  $\mathbf{w}$ ;

**Algorithm 1:** Active Sampling

enhances the efficiency and the performance. The iterative procedure continues until the budget  $B$  of chosen instances is fully consumed.

Another main difference between the active sampling scheme and typical pool-based active learning is that we sample  $b$  instances before the training step, while pool-based active learning often considers executing the training step right after querying the label of one instance. The difference is due to the fact that the pair-wise labels  $y_{ij}$  can be obtained easily and thus sampling and labeling can be relatively cheaper than querying in active learning. Furthermore, updating the weights right after knowing one instance may not lead to much improvement and can be too time consuming on large-scale data sets.

### 3.3. Sampling Strategies

Next, we discuss some possible sampling strategies that can be used in Algorithm 1. Besides a naïve choices that passively choose a random sample within  $\mathcal{U}^*$ , we define two measures that estimate the (learning) value of an unchosen pair. The two measures correspond to well-known criteria in pool-based active learning. Let  $\mathbf{x}_{ij}$  be the unchosen pair in  $\mathcal{U}^*$  with  $y_{ij} = 1$ , the two measures with respect to the current ranking function  $\mathbf{w}$  are

$$\text{closeness}(\mathbf{x}_{ij}, \mathbf{w}) = |\mathbf{w}^T \mathbf{x}_{ij}| \quad (3)$$

$$\text{correctness}(\mathbf{x}_{ij}, \mathbf{w}) = -[1 - \mathbf{w}^T \mathbf{x}_{ij}]_+ \quad (4)$$

The *closeness* measure corresponds to one of the most popular criteria in pool-based active learning called *uncertainty sampling* [23]. It captures the uncertainty of the ranking function  $\mathbf{w}$  on the unchosen pair. Intuitively, a low value of *closeness* means that the ranking function finds it hard to distinguish the two instances in the pair, which implies that the ranking function is less confident on the pair. Therefore, sampling the unchosen pairs that come with the lowest *closeness* values may improve the ranking performance by resolving the uncertainty.

On the other hand, the *correctness* measure is related to another common criterion in pool-based active learning called *expected error reduction* [25]. It captures the performance of the ranking function  $\mathbf{w}$  on the unchosen pair. Note that this exact correctness measure is only available within our active sampling scheme because we know the pair-label  $y_{ij}$  to always be 1 without loss of generality, while usual active learning algorithms do not know



the exact measure before querying and hence have to estimate it [10; 11]. A low value of *correctness* indicates that the ranking function does not perform well on the pair. Then, sampling the unchosen pairs that come with the lowest *correctness* values may improve the ranking performance by correcting the possible mistakes.

Similar to other active learning algorithms [10; 11], computing the pairs that come with the lowest *closeness* or *correctness* values can be time consuming, as it requires at least evaluating the values of  $\mathbf{w}^T \mathbf{x}_k$  for each instance  $(\mathbf{x}_k, y_k) \in \mathcal{D}$ , and then computing the measures on the pairs along with some selection or sorting steps that may be of super-linear time complexity [20]. Thus, such a *hard version* of active sampling is not computationally feasible for large-scale bipartite ranking. Next, we discuss the *soft version* of active sampling that randomly chooses pairs that come with lower *closeness* or *correctness* values by rejection sampling.

We consider a reject sampling step that samples a pair  $\mathbf{x}_{ij}$  with probability  $p_{ij}$ . We calculate a probability threshold to decide whether to accept the candidate pair. A pair that comes with a lower *closeness* or *correctness* values would enjoy a higher probability  $p_{ij}$  of being accepted.

Next, we define the probability threshold functions that correspond to the hard versions of *closeness* and *correctness*. Both threshold functions are in the shape of the sigmoid function, which is widely used to represent probabilities in logistic regression and neural networks [4]. For soft *closeness* sampling, we define the probability threshold  $p_{ij} \equiv 2 / (1 + e^{|\mathbf{w}^T \mathbf{x}_{ij}|})$ . For soft *correctness* sampling, we define  $p_{ij} \equiv 1 - 2 / (1 + e^{[1 - \mathbf{w}^T(\mathbf{x}_{ij})]_+})$ . We take different forms of soft versions because *closeness* is of range  $[0, \infty)$  while *correctness* is of range  $(-\infty, 0]$ .

Note that the sampling strategies above, albeit focusing on the most valuable pairs, is inheritedly biased. The chosen pool may not be representative enough of the whole training set because of the biased sampling strategies. There is a simple way that allows us to correct the sampling bias for learning a ranking function that performs well on the original bipartite ranking loss of interest. We take the idea of [19] to weight the sampled pair by the inverse of its probability of being sampled. That is, we could multiply the weight  $C_{ij}$  for a chosen pair  $\mathbf{x}_{ij}$  by  $\frac{1}{p_{ij}}$  when it gets returned by the rejection sampling.

### 3.4. Combined Ranking and Classification

Inspired by Theorem 1, the points can also carry some information for ranking. Next, we study how we can take those points into account during Active Sampling. We start by taking a closer look at the similarity and difference between the point-wise SVM (2) and the pair-wise SVM (1). The pair-wise SVM considers the weighted hinge loss on the pairs  $\mathbf{x}_{ij} = \mathbf{x}_i - \mathbf{x}_j$ , while the point-wise SVM considers the weighted hinge loss on the points  $\mathbf{x}_k$ . Consider one positive point  $(\mathbf{x}_k, +1)$ . Its hinge loss is  $[1 - \mathbf{w}^T \mathbf{x}_k]_+$ , which is the same as  $[1 - \mathbf{w}^T(\mathbf{x}_k - \mathbf{0})]_+$ . In other words, the positive point  $(\mathbf{x}_k, +1)$  can also be viewed as a *pseudo-pair* that consists of  $(\mathbf{x}_k, +1)$  and  $(\mathbf{0}, -1)$ . Similarly, a negative point  $(\mathbf{x}_j, -1)$  can be viewed as a pseudo-pair that consists of  $(\mathbf{x}_j, -1)$  and  $(\mathbf{0}, +1)$ . Let the set of all

pseudo-pairs within  $\mathcal{D}$  be

$$\begin{aligned} \mathcal{D}_{pseu} &= \{(\mathbf{x}_{i0} = \mathbf{x}_i - \mathbf{0}, +1) | \mathbf{x}_i \in \mathcal{D}^+\} \cup \{(\mathbf{x}_{0j} = \mathbf{0} - \mathbf{x}_j, +1) | \mathbf{x}_j \in \mathcal{D}^-\} \\ &\cup \{(\mathbf{x}_{0i} = \mathbf{0} - \mathbf{x}_i, -1) | \mathbf{x}_i \in \mathcal{D}^+\} \cup \{(\mathbf{x}_{j0} = \mathbf{x}_j - \mathbf{0}, -1) | \mathbf{x}_j \in \mathcal{D}^-\}. \end{aligned}$$

Then, the point-wise SVM (2) is just a variant of the pair-wise one (1) using the pseudo-pairs and some particular weights. Thus, we can easily unify the point-wise and the pair-wise SVMs together by minimizing some weighted hinge loss on the joint set  $\mathcal{D}^* = \mathcal{D}_{pair} \cup \mathcal{D}_{pseu}$  of pairs and pseudo-pairs. By introducing a parameter  $\gamma \in [0, 1]$  to control the relative importance between the real pairs and the pseudo-pairs, we propose the following novel formulation.

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^T \mathbf{w} + \gamma \sum_{\mathbf{x}_{ij} \in \mathcal{D}_{pair}^+} C_{crc}^{(ij)} [1 - \mathbf{w}^T \mathbf{x}_{ij}]_+ + (1 - \gamma) \sum_{\mathbf{x}_{k\ell} \in \mathcal{D}_{pseu}^+} C_{crc}^{(k\ell)} \cdot [1 - \mathbf{w}^T \mathbf{x}_{k\ell}]_+, \quad (5)$$

where  $\mathcal{D}_{pair}^+$  and  $\mathcal{D}_{pseu}^+$  denote the set of positive pairs and positive pseudo-pairs, respectively. The new formulation (5) combines the point-wise SVM and the pair-wise SVM in its objective function, and hence is named the Combined Ranking and Classification (CRC) framework. When  $\gamma = 1$ , CRC takes the pair-wise SVM (1) as a special case with  $C_{ij} = 2C_{crc}^{(ij)}$ ; when  $\gamma = 0$ , CRC takes the point-wise SVM (2) as a special case with  $C_+ = C_{crc}^{(i0)}$  and  $C_- = C_{crc}^{(0j)}$ . The CRC framework (5) remains as challenging to solve as the pair-wise SVM approach (1) because of the huge number of pairs. However, the general framework can be easily extended to the active sampling scheme, and hence be solved efficiently. We only need to change the training set from  $\mathcal{D}_{pair}$  to the joint set  $\mathcal{D}^*$ , and multiply the probability threshold  $p_{ij}$  in the soft version sampling by  $\gamma(1 - \gamma)$  for actual pairs (pseudo-pairs).

The CRC framework is closely related to the algorithm of Combined Ranking and Regression (CRR) [26] for general ranking. The CRR algorithm similarly considers a combined objective function of the point-wise terms and the pair-wise terms for improving the ranking performance. The main difference between CRR and CRC is that the CRR approach takes the squared loss on the points, while CRC takes the nature of bipartite ranking into account and considers the hinge loss on the points. The idea of combining pair-wise and point-wise approaches had been used in the multi-label classification problem [30]. The algorithm of Calibrated Ranking by Pairwise Comparison [17] assumes a calibration label between relevant and irrelevant labels, and unifies the pair-wise and point-wise label learning for multi-label classification.

To the best of our knowledge, while the CRR approach has reached promising performance in practice [26], the CRC formulation has not been seriously studied. The hinge loss used in CRC allows unifying the point-wise SVM and the pair-wise SVM under the same framework, and the unification is essential for applying *one* active sampling strategy on *both* the real pairs and the pseudo-pairs.

## 4. Experiments

We study the performance and efficiency of our proposed ASCRC algorithm on real-world large-scale data sets. We compare ASCRC with random-CRC, which does random sampling

Table 1: Data Sets Statistics

Dataset	positive	negative	total points	total pairs	dimension	AUC
letter*	789	19211	20000	30314958	16	CV
protein*	8198	9568	17766	156876928	357	test
news20	9999	9997	19996	199920006	1355191	CV
rcv1	10491	9751	20242	204595482	47236	CV
a9a	7841	24720	32561	387659040	123	test
bank	5289	39922	45211	422294916	51	CV
ijcnn1	4853	45137	49990	438099722	22	CV
shuttle*	34108	9392	43500	640684672	9	test
mnist*	5923	54077	60000	640596142	780	test
connect*	44473	23084	67557	2053229464	126	CV
acoustic*	18261	60562	78823	2211845364	50	test
real-sim	22238	50071	72309	2226957796	20958	CV
covtype	297711	283301	581012	168683648022	54	CV
url	792145	1603985	2396130	2541177395650	3231961	CV

under the CRC framework. In addition, we compare ASCRC with three other state-of-the-art algorithms: the point-wise weighted linear SVM (2) (WSVM), an efficient implementation [20] of the pair-wise linear RankSVM (1) (ERankSVM), and the combined ranking and regression (CRR) [26] algorithm for general ranking.

#### 4.1. Data Sets

We use 14 data sets from the LIBSVM Tools [8] and the UCI Repository [21] in the experiments. Table 1 shows the statistics of the data sets, which contains more than ten-thousands of instances and more than ten-millions of pairs. The data sets are definitely too large for a naïve implementation of RankSVM (1). Note that the data sets marked with (\*) are originally multi-class data sets, and we take the sub-problem of ranking the first class ahead of the other classes as a bipartite ranking task. For data sets that come with a moderate-sized test set, we report the test AUC. Otherwise we perform a 5-fold cross validation and report the cross-validation AUC.

#### 4.2. Experiment Settings

Given a budget  $B$  on the number of pairs to be used in each algorithm and a global regularization parameter  $C$ , we set the instance weights for each algorithm to fairly maintain the numerical scale between the regularization term and the loss terms. The global regularization parameter  $C$  is fixed to 0.1 in all the experiments. In particular, the setting below ensures that the total  $C_{(ij)}$ , summed over all the pairs (or pseudo-pairs), would be  $C \cdot B$  for all the algorithms.

- WSVM: As discussed in Section 2,  $C_+$  and  $C_-$  shall be inverse-proportional to  $N^+$  and  $N^-$  to make the weighted point-wise SVM a reasonable baseline for bipartite ranking. Thus, we set  $C_+ = \frac{B}{2N^+} \cdot C$  and  $C_- = \frac{B}{2N^-} \cdot C$  in (2). We solve the weighted SVM by the LIBLINEAR [13] package with its extension on instance weights.
- ERankSVM: We use the *SVM<sup>perf</sup>* [20] package to efficiently solve the linear RankSVM (1) with the AUC optimization option. We set the regularization parameter  $C_{perf} = \frac{B}{100} \cdot C$  where the 100 comes from a suggested value of the *SVM<sup>perf</sup>* package.
- CRR: We use the package *sofia-ml* [26] with the *sgd-svm* learner type, *combined-ranking* loop type and the default number of iterations that SGD takes to solve the problem. We set its regularization parameter  $\lambda = \frac{1}{B \cdot C}$ .

- ASCRC (ASRankSVM): We initialize  $|\mathcal{L}^*|$  to  $b$ , and assign  $C_{crc}^{(ij)} = \frac{\Gamma|\mathcal{L}^*|}{p_{ij}Z} \cdot C$  in each iteration, where  $\Gamma$  equals to either  $\gamma$  or  $(1-\gamma)$  for either real or pseudo pairs and  $Z$  is a normalization constant  $\sum_{\mathbf{x}_{ij} \in \mathcal{L}^*} \frac{1}{p_{ij}}$  that prevents  $C_{crc}^{(ij)}$  from being too large. We solve the linearSVM within ASCRC by the LIBLINEAR [13] package with its extension on instance weights.
  - random-CRC: random-CRC corresponds to ASCRC with  $p_{ij} = 1$  for all the pairs.
- To evaluate the average performance of ASCRC and random-CRC algorithms, we average their results over 10 different initial pools.

### 4.3. Performance Comparison and Robustness

We first set  $\gamma = 1$  in ASCRC and random-CRC, which makes ASCRC equivalent to ASRankSVM. The more general ASCRC will be studied later. We let  $b = 100$  and  $B = 8000$ , which is a relatively small budget out of the millions of pairs.

We show how the AUC changes as  $|\mathcal{L}^*|$  grows throughout the active sampling steps of ASRankSVM in Fig. 1. For WSVM, ERankSVM and CRR, we plot a horizontal line on the AUC achieved when using the whole training set. We also list the final AUC with the standard deviation of all the algorithms in Table 2.

From Fig. 1 and Table 2, we see that soft-correct are generally the best among the other competitors. Further, we conduct right-tail t-test for soft-correct against the others to show whether the improvement of soft-correct sampling is significant. In Table 3, we summarize the performance result under a 95% significance level.

First, we compare soft-correct with random sampling and discover that soft-correct performs better on 10 data sets, which shows that active sampling is working reasonably well. While comparing soft-close with soft-correct in Table 2 and Table 3, we find that soft-correct outperforms soft-close on 7 data sets and ties with 5. Moreover, Fig. 1 shows the strong performance of soft-correct comes from the early steps of active sampling. Finally, when comparing soft-correct with other algorithms, we discover that soft-correct performs the best on 8 data sets: it outperforms ERankSVM on 8 data sets, WSVM on 9 data sets, and CRR on 11 data sets. The results demonstrate that even when using a pretty small sampling budget of 8,000 pairs, ASRankSVM with soft-correct sampling can achieve significant improvement over those state-of-the-art ranking algorithms that use the whole

Figure 1: Performance Curves on Different Datasets

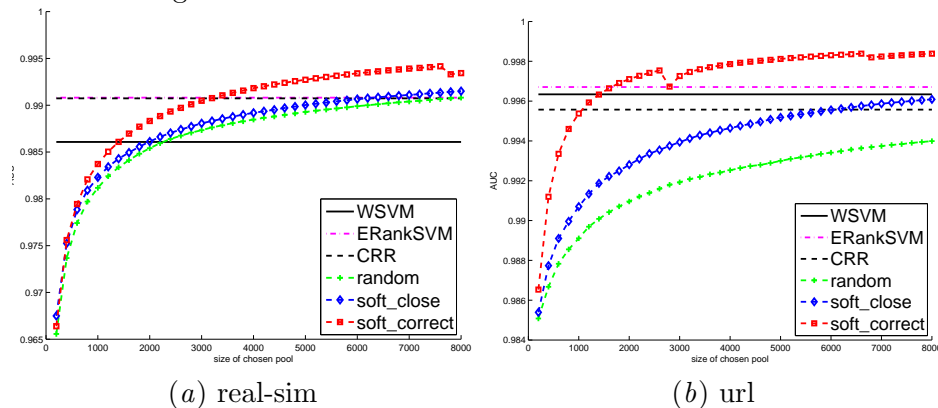


Table 2: AUC (mean $\pm$ std) at  $|\mathcal{L}^*| = 8000, b = 100, \gamma = 1.0$ 

Data	WSVM	ERankSVM	CRR	Random	ASRankSVM	
					Soft-Close	Soft-Correct
letter	.9808	.9877	.9874	.9883 $\pm$ .0003	<b>.9883</b> $\pm$ .0002	.9874 $\pm$ .0123
protein	<b>.8329</b>	.8302	.8306	.8229 $\pm$ .0031	.8240 $\pm$ .0016	.8233 $\pm$ .0028
news20	.9379	.9753	.9743	.9828 $\pm$ .0008	.9836 $\pm$ .0006	<b>.9903</b> $\pm$ .0003
rcv1	.9876	.9916	.9755	.9920 $\pm$ .0004	.9923 $\pm$ .0003	<b>.9944</b> $\pm$ .0002
a9a	.9008	<b>.9047</b>	.8999	.9003 $\pm$ .0006	.9012 $\pm$ .0004	.9007 $\pm$ .0006
bank	.8932	.9023	.8972	.9051 $\pm$ .0010	.9057 $\pm$ .0011	<b>.9083</b> $\pm$ .0007
ijcnn1	.9335	.9343	.9336	.9342 $\pm$ .0004	.9345 $\pm$ .0006	<b>.9348</b> $\pm$ .0003
shuttle	.9873	.9876	.9888	.9894 $\pm$ .0001	.9896 $\pm$ .0001	<b>.9907</b> $\pm$ .0000
mnist	<b>.9985</b>	.9983	.9973	.9967 $\pm$ .0004	.9979 $\pm$ .0001	.9976 $\pm$ .0002
connect	.8603	<b>.8613</b>	.8532	.8594 $\pm$ .0008	.8604 $\pm$ .0007	.8603 $\pm$ .0009
acoustic	.8881	.8911	.8931	.8952 $\pm$ .0005	.8952 $\pm$ .0005	<b>.8988</b> $\pm$ .0004
real-sim	.9861	.9908	.9907	.9908 $\pm$ .0003	.9915 $\pm$ .0002	<b>.9934</b> $\pm$ .0064
covtype	.8047	.8228	.8189	.8238 $\pm$ .0008	.8239 $\pm$ .0007	<b>.8249</b> $\pm$ .0006
url	.9963	.9967	.9956	.9940 $\pm$ .0003	.9961 $\pm$ .0001	<b>.9984</b> $\pm$ .0015

training data set. Also, the tiny standard deviation shown in Table 2 and the significant results from t-test suggest the robustness of ASRankSVM with soft-correct in general.

#### 4.4. Efficiency Comparison

First, we study the efficiency of soft active sampling by checking the average number of rejected samples before passing the probability threshold during rejection sampling. The number is plotted against the size of  $\mathcal{L}^*$  in Fig. 2. The soft-close strategy usually needs fewer than 10 rejected samples, while the soft-correct strategy generally needs an increasing number of rejected samples. The reason is that when the ranking performance becomes better throughout the iterations, the probability threshold behind soft-correct could be pretty small. The results suggest that the soft-close strategy is generally efficient, while the soft-correct strategy may be less efficient as  $\mathcal{L}^*$  grows.

Next, we list the CPU time consumed for all algorithms under 8,000 pairs budget in Table 4, and the data sets are ordered ascendantly by its size. We can see that WSVM and CRR run fast but give inferior performance while ERankSVM performs better but the training time grows fast as the data size increases. The result is consistent with the discussion in Section 1 that conducting bipartite ranking efficiently and accurately at the same time is challenging. On the other hand, under the ASRankSVM scheme, random runs the fastest, then soft-close, and soft-correct is the slowest. This result reflects the average number of rejected samples discussed above. Also, we can find that the soft version samplings are usually much faster than the corresponding hard versions, which validate the time consuming enumerating or sorting steps can not fit our goal in terms of efficiency.

More importantly, when comparing soft-correct with ERankSVM, soft-correct runs faster on 7 data sets, which suggests the development of ASRankSVM is as efficient as the state-

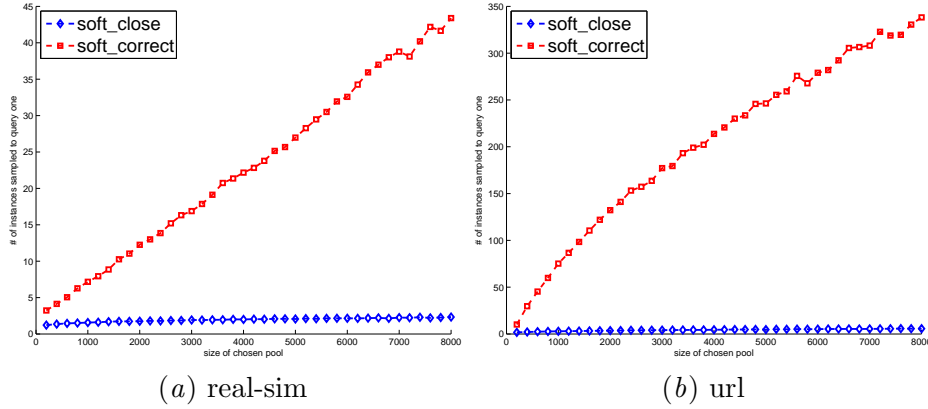
Table 3: 95% t-test Results Summary: Soft-Correct Versus Other Algorithms at  $|\mathcal{L}^*| = 8000, b = 100, \gamma = 1.0$ 

	WSVM	ERankSVM	CRR	Random	Soft-Close
Total (win/loss/tie)	9/2/3	8/4/2	11/1/2	10/0/4	7/2/5

Table 4: CPU Time Table under 8000 Pair Budget(Seconds)

Data	Random	Soft-Close	Soft-Correct	Hard-Close	Hard-Correct	WSVM	ERankSVM	CRR
letter	0.808	1.53	19.602	192.671	42.553	0.29	0.72	0.15
protein	3.136	2.943	4.29	12.128	8.315	0.85	4.9	0.99
news20	58.594	56.506	66.394	184.233	128.056	20.1	10.64	4.55
rcv1	5.872	6.318	21.789	114.028	35.258	1.83	2.54	0.77
a9a	0.374	0.504	1.065	30.384	19.537	0.28	4.25	0.31
bank	1.957	2.301	4.644	20.8064	13.8512	0.142	5.624	0.3
ijcnn1	0.957	1.508	4.031	107.002	79.32	0.69	2.75	0.28
shuttle	0.146	0.288	3.02	26.307	17.577	0.18	0.98	0.36
mnist	1.61	2.851	56.604	205.135	50.174	4.92	22.38	2.75
connect	2.734	3.229	5.493	117.047	121.359	2.5	15.42	0.78
acoustic	0.488	0.624	1.03	33.39	41.167	1.82	11.57	1.93
real-sim	4.025	4.35	11.648	318.702	139.042	3.15	7.58	1.78
covtype	1.47	1.801	2.739	800.539	5739.97	6.41	29.54	3.3
url	31.026	31.011	163.022	122045.33	285394.18	116.61	594.82	58.1

Figure 2: Sampling Efficiency Curves on Different Datasets



of-the-art ERankSVM on large-scale data sets in general. Nevertheless, we can find that the CPU time of soft-correct grows much slower than ERankSVM as data size increases because the time complexity of ASRankSVM scheme mainly depends on the budget  $B$  and the step size  $b$ , not the size of data.

Figure 3: The Benefit of Trade-off Parameter

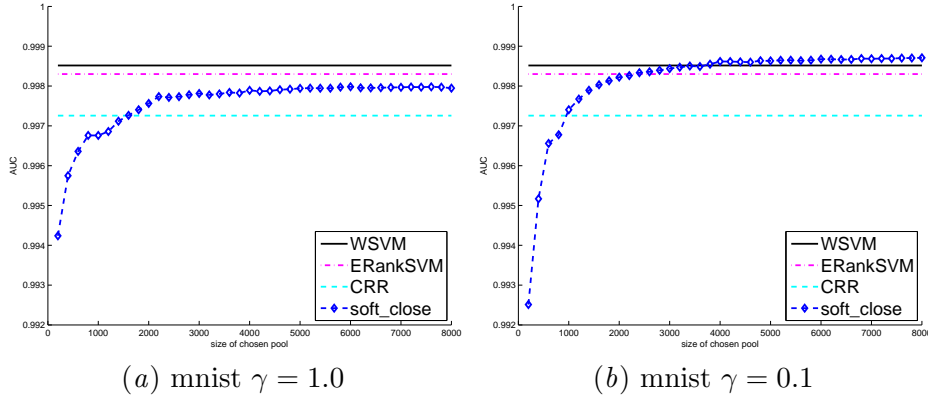


Table 5: Optimal  $\gamma$  on Different Datasets

	letter	protein	news20	rcv1	a9a	bank	ijcnn1
Soft-Close	<b>uniform</b>	1	<b>uniform,1</b>	<b>uniform,1</b>	uniform,0.7	<b>1</b>	1
Soft-Correct	<b>0.9</b>	0.1,0.2	<b>uniform,1</b>	<b>uniform,1</b>	uniform,0.6,0.7	<b>uniform,1</b>	<b>0.8,0.9,1</b>
	shuttle	mnist	connect	acoustic	real-sim	covtype	url
Soft-Close	<b>uniform,1</b>	<b>0.1,0.2,0.4</b>	uniform,1	<b>1</b>	<b>uniform,1</b>	<b>uniform,1</b>	uniform,1
Soft-Correct	<b>uniform,1</b>	0.6	uniform,1	<b>uniform,1</b>	<b>uniform</b>	<b>uniform,0.9,1</b>	<b>0.8</b>

#### 4.5. The Usefulness of the CRC Framework

Next, we study the necessity of the CRC framework by comparing the performance of soft-closeness and soft-correctness under different choices of  $\gamma$ . We report the best  $\gamma$  within  $\{\textit{uniform}, 0.1, 0.2, \dots, 1.0\}$ , where *uniform* means balancing the influence of real pairs and pseudo-pairs by  $\gamma = \frac{|\mathcal{D}_{\textit{pair}}|}{|\mathcal{D}^*|}$ . Table 5 shows the best  $\gamma$  on each data set and sampling strategy, and the bold  $\gamma$  indicates the setting outperforms ERankSVM significantly. We see that using  $\gamma = 1$  (real pairs only) performs well in most data sets, while a smaller  $\gamma$  or *uniform* can sometimes reach the best performance. The results justify that the real pairs are more important than the pseudo-pairs, while the latter can sometimes be helpful. When pseudo-pair helps, as shown in Fig. 3 for the mnist data set, the flexibility of the CRC framework can be useful.

## 5. Conclusion

We propose the algorithm of Active Sampling (AS) under Combined Ranking and Classification (CRC) based on the linear SVM. There are two major components of the proposed algorithm. The AS scheme selects valuable pairs for training and resolves the computational burden in large-scale bipartite ranking. The CRC framework unifies the concept of point-wise ranking and pair-wise ranking under the same framework, and can perform better than pure point-wise ranking or pair-wise ranking. The unified view of pairs and points (pseudo-pairs) in CRC allows using one AS scheme to select from both types of pairs.

Experiments on 14 real-world large-scale data sets demonstrate the promising performance and efficiency of the ASRankSVM and ASCRC algorithms. The algorithms usually outperform state-of-the-art bipartite ranking algorithms, including the point-wise SVM, the pair-wise SVM, and the combined ranking and regression approach. The results not only justify the validity of ASCRC, but also shows the valuable pairs or pseudo-pairs can be helpful for large-scale bipartite ranking.

## Acknowledgments

We thank Prof. Shou-De Lin, Yuh-Jye Lee, the anonymous reviewers and the members of the NTU Computational Learning Lab for valuable suggestions. This work is part of the first author’s Master thesis [28] and is supported by the National Science Council of Taiwan via the grant NSC 101-2628-E-002-029-MY2.

## References

- [1] S. Agarwal and D. Roth. *A study of the bipartite ranking problem in machine learning*. University of Illinois at Urbana-Champaign, 2005.

- [2] N. Ailon. An active learning algorithm for ranking from pairwise preferences with an almost optimal query complexity. *JMLR*, 13:137–164, 2012.
- [3] M.-F. Balcan, N. Bansal, A. Beygelzimer, D. Coppersmith, J. Langford, and G. B. Sorkin. Robust reductions from ranking to classification. *Machine learning*, 72(1-2):139–153, 2008.
- [4] E. B. Baum and F. Wilczek. Supervised learning of probability distributions by neural networks. In *NIPS*, pages 52–61, 1988.
- [5] U. Brefeld and T. Scheffer. AUC maximizing support vector learning. In *ICML Workshop on ROC Analysis in Machine Learning*, 2005.
- [6] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender. Learning to rank using gradient descent. In *ICML*, pages 89–96, 2005.
- [7] R. Caruana, T. Joachims, and L. Backstrom. KDD-Cup 2004: results and analysis. *ACM SIGKDD Explorations Newsletter*, 6(2):95–108, 2004.
- [8] C.-C. Chang and C.-J. Lin. LIBSVM: A library for support vector machines. *ACM TIST*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- [9] C. Cortes and M. Mohri. AUC optimization vs. error rate minimization. *NIPS*, 16(16):313–320, 2004.
- [10] P. Donmez and J. G. Carbonell. Optimizing estimated loss reduction for active sampling in rank learning. In *ICML*, pages 248–255, 2008.
- [11] P. Donmez and J. G. Carbonell. Active sampling for rank learning via optimizing the area under the ROC curve. *AIR*, pages 78–89, 2009.
- [12] Ş. Ertekin and C. Rudin. On equivalence relationships between classification and ranking algorithms. *JMLR*, 12:2905–2929, 2011.
- [13] R.-E. Fan, K.-W. Chang, C.-J. Hsieh, X.-R. Wang, and C.-J. Lin. LIBLINEAR: A library for large linear classification. *JMLR*, 9:1871–1874, 2008.
- [14] T. Fawcett. An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8):861–874, 2006.
- [15] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer. An efficient boosting algorithm for combining preferences. *JMLR*, 4:933–969, 2003.
- [16] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *JCSS*, 55(1):119–139, 1997.
- [17] J. Fürnkranz, E. Hüllermeier, E. L. Mencia, and K. Brinker. Multilabel classification via calibrated label ranking. *Machine Learning*, 73(2):133–153, 2008.
- [18] R. Herbrich, T. Graepel, and K. Obermayer. Large margin rank boundaries for ordinal regression. *NIPS*, pages 115–132, 1999.
- [19] D. G. Horvitz and D. J. Thompson. A generalization of sampling without replacement from a finite universe. *JASA*, 47(260):663–685, 1952.
- [20] T. Joachims. Training linear svms in linear time. In *ACM SIGKDD*, pages 217–226, 2006.
- [21] M. L. Kevin Bache. UCI machine learning repository, 2013.
- [22] W. Kotłowski, K. Dembczynski, and E. Hüllermeier. Bipartite ranking through minimization of univariate loss. In *ICML*, pages 1113–1120, 2011.
- [23] D. D. Lewis and W. A. Gale. A sequential algorithm for training text classifiers. In *ACM SIGIR*, pages 3–12, 1994.
- [24] T.-Y. Liu. Learning to rank for information retrieval. *FTIR*, 3(3):225–331, 2009.
- [25] N. Roy and A. McCallum. Toward optimal active learning through monte carlo estimation of error reduction. In *ICML*, pages 441–448, 2001.
- [26] D. Sculley. Combined regression and ranking. In *ACM SIGKDD*, pages 979–988, 2010.
- [27] B. Settles. Active learning literature survey. *University of Wisconsin, Madison*, 2010.
- [28] W.-Y. Shen and H.-T. Lin. Active sampling of pairs and points for large-scale linear bipartite ranking. Master’s thesis, National Taiwan University, 2013.
- [29] H. Steck. Hinge rank loss and the area under the ROC curve. *ECML*, pages 347–358, 2007.
- [30] G. Tsoumakas and I. Katakis. Multi-label classification: An overview. *IJDWM*, 3(3):1–13, 2007.
- [31] V. Vapnik. *The nature of statistical learning theory*. Springer, 1999.
- [32] H. Yu. SVM selective sampling for ranking with application to data retrieval. In *ACM SIGKDD*, pages 354–363, 2005.
- [33] G.-X. Yuan, C.-H. Ho, and C.-J. Lin. Recent advances of large-scale linear classification. *Proceedings of the IEEE*, 100(9):2584–2603, 2012.