

Active Learning for Ranking through Expected Loss Optimization

Bo Long, Jiang Bian, Olivier Chapelle, Ya Zhang, Yoshiyuki Inagaki, and Yi Chang

Abstract—Learning to rank arises in many data mining applications, ranging from web search engine, online advertising to recommendation system. In learning to rank, the performance of a ranking model is strongly affected by the number of labeled examples in the training set; on the other hand, obtaining labeled examples for training data is very expensive and time-consuming. This presents a great need for the active learning approaches to select most informative examples for ranking learning; however, in the literature there is still very limited work to address active learning for ranking. In this paper, we propose a general active learning framework, expected loss optimization (ELO), for ranking. The ELO framework is applicable to a wide range of ranking functions. Under this framework, we derive a novel algorithm, expected discounted cumulative gain (DCG) loss optimization (ELO-DCG), to select most informative examples. Then, we investigate both query and document level active learning for raking and propose a two-stage ELO-DCG algorithm which incorporate both query and document selection into active learning. Furthermore, we show that it is flexible for the algorithm to deal with the skewed grade distribution problem with the modification of the loss function. Extensive experiments on real-world web search data sets have demonstrated great potential and effectiveness of the proposed framework and algorithms.

Index Terms—Active learning, ranking, expected loss optimization

1 INTRODUCTION

RANKING is the core component of many important information retrieval problems, such as web search, recommendation, computational advertising. Learning to rank represents an important class of supervised machine learning tasks with the goal of automatically constructing ranking functions from training data. As many other supervised machine learning problems, the quality of a ranking function is highly correlated with the amount of labeled data used to train the function. Due to the complexity of many ranking problems, a large amount of labeled training examples is usually required to learn a high quality ranking function. However, in most applications, while it is easy to collect unlabeled samples, it is very expensive and time-consuming to label the samples.

Active learning comes as a paradigm to reduce the labeling effort in supervised learning. It has been mostly studied in the context of classification tasks [20]. Existing algorithms for learning to rank may be categorized into three groups: pointwise approach [8], pairwise approach [26], and listwise

approach [22]. Compared to active learning for classification, active learning for ranking faces some unique challenges. First, there is no notion of classification margin in ranking. Hence, many of the margin-based active learning algorithms proposed for classification tasks are not readily applicable to ranking. Further more, even some straightforward active learning approach, such as query-by-committee (QBC), has not been justified for the ranking tasks under regression framework. Second, in most supervised learning setting, each data sample can be treated completely independent of each other. In learning to rank, data examples are not independent, though they are conditionally independent given a query. We need to consider this data dependence in selecting data and tailor active learning algorithms according to the underlying learning to rank schemes. Third, ranking problems are often associated with very skewed data distributions. For example, in the case of document retrieval, the number of irrelevant documents is of orders of magnitude more than that of relevant documents in training data. It is desirable to consider the data skewness when selecting data for ranking. Therefore, there is a great need for an active learning framework for ranking.

In this paper, we attempt to address those three important and challenging aspects of active learning for ranking. We first propose a general active learning framework, expected loss optimization (ELO), and apply it to ranking. The key idea of the proposed framework is that given a loss function, the samples minimizing the expected loss (EL) are the most informative ones. Under this framework, we derive a novel active learning algorithm for ranking, which uses function ensemble to select most informative examples that minimizes a chosen loss. For the rest of the paper, we use web search ranking problem as an example to illustrate

- B. Long is with LinkedIn Inc., Mountain View, CA 94043. E-mail: blong@linkedin.com.
- J. Bian is with Microsoft Research, Beijing, China 100080. E-mail: jibian@microsoft.com.
- O. Chapelle is with Criteo, Palo Alto, CA 94301. E-mail: olivier@chapelle.cc.
- Y. Zhang is with Shanghai Jiao Tong University, Shanghai, China 200240. E-mail: ya_zhang@sjtu.edu.cn.
- Y. Inagaki and Y. Chang are with Yahoo! Labs, Sunnyvale, CA 94089. E-mail: {inagaki, yichang}@yahoo-inc.com.

Manuscript received 3 May 2013; revised 22 Dec. 2013; accepted 13 Sept. 2014. Date of publication 29 Oct. 2014; date of current version 27 Mar. 2015.

Recommended for acceptance by I. Davidson.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TKDE.2014.2365785

the ideas and perform evaluation. But the proposed method is generic and may be applicable to all ranking applications. In the case of web search ranking, we minimize the expected discounted cumulative gain (DCG) loss, one of the most commonly used loss for web search ranking. This algorithm may be easily adapted to other ranking loss such as normalized discounted cumulative gain (NDCG) or average precision. To address the data dependency issue, the proposed algorithm is further extended to a two-stage active learning schema to seamlessly integrate query level and document level data selection. Finally, we extended the proposed algorithms to address the skew grade distribution problem.

The main contributions of the paper are summarized as follows.

- We propose a general active learning framework based on expected loss optimization. This framework is applicable to various ranking scenarios with a wide spectrum of learners and loss functions. We also provides a theoretically justified measure of the informativeness.
- Under the ELO framework, we derive novel algorithms to select most informative examples by optimizing the expected DCG loss. Those selected examples represent the ones that the current ranking model is most uncertain about and they may lead to a large DCG loss if predicted incorrectly.
- We propose a two stage active learning algorithm for ranking, which addresses the sample dependence issue by first performing query level selection and then document level selection.
- We further show how to extend ELO framework to address the skewed grade distribution problem in ranking. Balanced version ELO algorithms are derived for both query level active learning and two-stage active learning.

2 RELATED WORK

The main motivation for active learning is that it usually requires time and/or money for the human expert to label examples and those resources should not be wasted to label non-informative samples, but be spent on interesting ones.

Optimal experimental design [12] is closely related to active learning as it attempts to find a set of points such that the variance of the estimate is minimized. In contrast to this “batch” formulation, the term *active learning* often refers to an incremental strategy [7].

There has been various types of strategies for active learning that we now review. A comprehensive survey can be found in [21]. The simplest and maybe most common strategy is *uncertainty sampling* [18], where the active learning algorithm queries points for which the label uncertainty is the highest. The drawback of this type of approach is that it often mixes two types of uncertainties, the one stemming from the *noise* and the *variance*. The noise is something intrinsic to the learning problem which does not depend on the size of the training set. An active learner should not spend too much effort in querying points in noisy regions of the input space. On the other hand, the variance is the

uncertainty in the model parameters resulting from the finiteness of the training set. Active learning should thus try to minimize this variance and this was first proposed in [7].

In Bayesian terms, the variance is computed by integrating over the posterior distribution of the model parameters. But in practice, it may be difficult or impossible to compute this posterior distribution. Instead, one can randomly sample models from the posterior distribution [9]. An heuristic way of doing so is to use a bagging type of algorithm [1]. This type of analysis can be seen as an extension of the *Query-By-Committee* algorithm [14] which has been derived in a noise free classification setting. In that case, the posterior distribution is uniform over the *version space*—the space of consistent hypothesis with the labeled data—and the QBC algorithm selects points on which random functions in the version space have the highest disagreement.

Another fairly common heuristic for active learning is to select points that once added in the training set are expected to result in a large model change [21] or a large increase in the objective function value that is being optimized [4].

Another related field is ranking. In recent years, the ranking problem is frequently formulated as a supervised machine learning problem [28], [29], [30], [31], [32], [33], [34]. These learning-to-rank approaches are capable of combining different kinds of features to train ranking functions. The problem of ranking can be formulated as that of learning a ranking function from pair-wise preference data. The idea is to minimize the number of contradicting pairs in training data. For example, RankSVM [28] uses support vector machines to learn a ranking function from preference data. RankNet [29] applies neural network and gradient descent to obtain a ranking function. RankBoost [30] applies the idea of boosting to construct an efficient ranking function from a set of weak ranking functions. The studies reported in [33] proposed a framework called GBRank using gradient descent in function spaces, which is able to learn relative ranking information in the context of web search.

For the main focus of this paper, active learning for ranking, compared with traditional active learning, there is still limited work in the literature. Donmez and Carbonell studied the problem of document selection in ranking [11]. Their algorithm selects the documents which, once added to the training set, are the most likely to result in a large change in the model parameters of the ranking function. They apply their algorithm to RankSVM [17] and RankBoost [13]. Also in the context of RankSVM, [25] suggests to add the most ambiguous pairs of documents to the training set, that is documents whose predicted relevance scores are very close under the current model [35], [36]. Other works based on pairwise ranking include [6], [10]. In case of binary relevance, [5] proposed a greedy algorithm which selects document that are the most likely to differentiate two ranking systems in terms of average precision. Moreover, an empirical comparison of document selection strategies for learning to rank can be found in [2]. Finally, Long et al. [19] proposed a general active learning framework for ranking. But, it still suffered from the problem of imbalanced grade in ranking.

There are some related works about query sampling. Yilmaz and Robertson [24] empirically show that having more queries but less number of documents per query is better than having more documents and less queries. Yang

et al. propose a greedy query selection algorithm that tries to maximize a linear combination of query difficulty, query density and query diversity [23].

3 EXPECTED LOSS OPTIMIZATION FOR ACTIVE LEARNING

As explained in the previous section, a natural strategy for active learning is based on variance minimization. The variance, in the context of regression, stems from the uncertainty in the prediction due to the finiteness of the training set. Cohn et al. [7] proposes to select the next instance to be labeled as the one with the highest variance. However, this approach applies only to regression and we aim at generalizing it through the *Bayesian expected loss* [3].

In the rest of the section, we first review Bayesian decision theory in Section 3.1 and then introduce the *expected loss optimization* principle for active learning. In Section 3.2 we show that in the cases of classification and regression, applying ELO turns out to be equivalent to standard active learning method. Finally, we present ELO for ranking in Section 3.3.

3.1 Bayesian Decision Theory

We consider a classical Bayesian framework to learn a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ parametrized by a vector θ . The training data D is made of n examples, $(x_1, y_1), \dots, (x_n, y_n)$. Bayesian learning consists in:

- 1) Specifying a prior $P(\theta)$ on the parameters and a likelihood function $P(y|x, \theta)$.
- 2) Computing the likelihood of the training data, $P(D|\theta) = \prod_{i=1}^n P(y_i|x_i, \theta)$.
- 3) Applying Bayes rule to get the *posterior* distribution of the model parameters, $P(\theta|D) = P(D|\theta)P(\theta)/P(D)$.
- 4) For a test point x , computing the *predictive distribution* $P(y|x, D) = \int_{\theta} P(y|x, \theta)P(\theta|D)d\theta$.

Note that in such a Bayesian formalism, the prediction is a distribution instead of an element of the output space \mathcal{Y} . In order to know which action to perform (or which element to predict), Bayesian decision theory needs a *loss* function. Let $\ell(a, y)$ be the loss incurred by performing action a when the true output is y . Then the Bayesian expected loss is defined as the expected loss under the predictive distribution:

$$\rho(a) := \int_{\mathcal{Y}} \ell(a, y)P(y|x, D)dy. \quad (1)$$

The best action according to Bayesian decision theory is the one that minimizes that loss: $a^* := \arg \min_a \rho(a)$. Central to our analysis is the expected loss of that action, $\rho(a^*)$ or

$$\text{EL}(x) := \min_a \int_{\mathcal{Y}} \ell(a, y)P(y|x, D)dy. \quad (2)$$

This quantity should be understood as follows: given that we have taken the best action a^* for the input x , and that the true output is in fact given by $P(y|x, D)$, what is, in expectation, the loss to be incurred once the true output is revealed?

The overall generalization error (i.e. the expected error on unseen examples) is the average of the expected loss over the input distribution: $\int_{\mathcal{X}} \text{EL}(x)P(x)dx$. Thus, in order to minimize this generalization error, our active learning strategy consists in selecting the input instance x to maximize the expected loss:

$$\arg \max_x \text{EL}(x).$$

3.2 ELO for Regression and Classification

In this section, we show that the ELO principle for active learning is equivalent to well known active learning strategies for classification and regression. In the cases of regression and classification, the "action" a discussed above is simply the prediction of an element in the output space \mathcal{Y} .

3.2.1 Regression

The output space for regression is $\mathcal{Y} = \mathbb{R}$ and the loss function is the squared loss $\ell(a, y) = (a - y)^2$. It is well known that the prediction minimizing this square loss is the mean of the distribution and that the expected loss is the variance:

$$\arg \min_a \int_{\mathcal{Y}} (a - y)^2 P(y|x, D)dy = \mu$$

and

$$\min_a \int_{\mathcal{Y}} (a - y)^2 P(y|x, D)dy = \sigma^2,$$

where μ and σ^2 are the mean and variance of the predictive distribution. So in the regression case, ELO will choose the point with the highest predictive variance which is exactly one of the classical strategy for active learning [7].

3.2.2 Classification

The output space for binary classification is $\mathcal{Y} = \{-1, 1\}$ and the loss is the 0/1 loss: $\ell(a, y) = 0$ if $a = y$, 1 otherwise. The optimal prediction is given according to $\arg \max_{a \in \mathcal{Y}} P(y = a|x, D)$ and the expected loss turns out to be:

$$\min(P(y = 1|x, D), P(y = -1|x, D)),$$

which is maximum when $P(y = 1|x, D) = P(y = -1|x, D) = 0.5$, that is when we are completely uncertain about the class label. This uncertainty based active learning is the most popular one for classification which was first proposed in [18].

3.3 ELO for Ranking

In the case of ranking, the input instance is a query and a set of documents associated with it, while the output is a vector of relevance scores. If the query q has n documents, let us denote by $X_q := (x_1, \dots, x_n)$ the feature vectors describing these (query, document) pairs and by $Y := (y_1, \dots, y_n)$ their labels. As before we have a predictive distribution $P(Y|X_q, D)$. Unlike active learning for classification and regression, active learning for ranking can select examples at different levels. One is query level, which selects a query with all associated documents X_q ; the other one is document level, which selects documents x_i individually.

3.3.1 Query Level

In the case of ranking, the "action" in ELO framework is slightly different than before because we are not directly interested in predicting the scores, but instead we want to produce a ranking. So the set of actions is the set of permutations of length n and for a given permutation π , the rank of the i th document $\pi(i)$. The expected loss for a given π can thus be written as:

$$\int_Y \ell(\pi, Y) P(Y|X_q, D) dY, \quad (3)$$

where $\ell(\pi, Y)$ quantifies the loss in ranking according to π if the true labels are given by Y . The next section will detail the computation of the expected loss where ℓ is the DCG loss.

As before, the ELO principle for active learning tells us to select the queries with the highest expected losses:

$$EL(q) := \min_{\pi} \int_Y \ell(\pi, Y) P(Y|X_q, D) dY. \quad (4)$$

As an aside, note that the ranking minimizing the loss (3) is not necessarily the one obtained by sorting the documents according to their mean predicted scores. This has already been noted for instance in [27].

3.3.2 Document Level

Selecting the most informative document is a bit more complex because the loss function in ranking is defined at the query level and not at the document level. We can still use the expected loss (4), but only consider the predictive distribution for the document of interest and consider the scores for the other documents fixed. Then we take an expectation over the scores of the other documents. This leads to:

$$EL(q, i) = \int_{Y^i} \min_{\pi} \int_{y_i} \ell(\pi, Y) P(Y|X_q, D) dy_i dY^i, \quad (5)$$

where $EL(q, i)$ is the expected loss for query q associated with the i th document and Y^i is the vector Y after removing y_i .

4 ALGORITHM DERIVATION

We now provide practical implementation details of the ELO principle for active learning and in particular specify how to compute equations (4) and (5) in case of the DCG loss.

The difficulty of implementing the formulations of the previous section lies in the fact that the computation of the posterior distributions $P(y_i|x_i, D)$ and the integrals is in general intractable. For this reason, we instead use an ensemble of learners and replace the integrals by sums over the ensemble. As in [1], we propose to use bootstrap to construct the ensemble. More precisely, the labeled set is subsampled several times and for each subsample, a relevance function is learned. The predictive distribution for a document is then given by the predicted relevance scores by various functions in the ensemble. The use of bootstrap to estimate predictive distributions is not new and there has been some work investigating whether the two procedures are equivalent [16].

Finally note that in our framework we need to estimate the relevance scores; therefore, we concentrate in this paper

on pointwise approaches for learning to rank, since pairwise and listwise approaches may not provide relevance scores. However, the proposed algorithm can be applied to situations when pairwise and listwise approaches can generate relevance scores.

4.1 Query Level Active Learning

If the metric of interest is DCG, the associated loss is the difference between the DCG for that ranking and the ranking with largest DCG:

$$\ell(\pi, Y) = \max_{\pi'} \text{DCG}(\pi', Y) - \text{DCG}(\pi, Y), \quad (6)$$

where

$$\text{DCG}(\pi, Y) = \sum_i \frac{2^{y_i} - 1}{\log_2(1 + \pi(i))}. \quad (7)$$

Substituting (6) into (4), the expected loss for a given q is expressed as follows:

$$\begin{aligned} EL(q) &= \min_{\pi} \int_Y (\max_{\pi'} \text{DCG}(\pi', Y) - \text{DCG}(\pi, Y)) \\ &\quad P(Y|X_q, D) dY \\ &= \int_Y \max_{\pi'} \text{DCG}(\pi', Y) P(Y|X_q, D) dY \\ &\quad - \max_{\pi} \int_Y \text{DCG}(\pi, Y) P(Y|X_q, D) dY. \end{aligned} \quad (8)$$

The derivation of the second step uses the fact that the first term in the loss function (6) is a constant w.r.t. π .

The maximum in the first component of (8) can easily be found by sorting the documents according to Y . We rewrite the integral in the second component as:

$$\begin{aligned} &\int_Y \text{DCG}(\pi, Y) P(Y|X_q, D) dY \\ &= \sum_i \frac{1}{\log_2(1 + \pi(i))} \underbrace{\int_{y_i} (2^{y_i} - 1) P(y_i|x_i, D) dy_i}_{:=t_i}, \end{aligned} \quad (9)$$

with which the maximum can now be found by sorting the t_i .

The pseudo-code for selecting queries based on equation (8) is presented in Algorithm 1. The notations and definitions are as follows:

- G is the *gain* function defined as

$$G(s) = 2^s - 1.$$

- The notation $\langle \cdot \rangle$ means average. For instance,

$$\langle d_i \rangle = \frac{1}{N} \sum_{i=1}^N d_i.$$

- BDCG is a function which takes as input a set of gain values and returns the corresponding best DCG:

$$BDCG(\{g_j\}) = \sum_j \frac{g_j}{\log_2(1 + \pi^*(j))},$$

where π^* is the permutation sorting the g_j in decreasing order.

Algorithm 1. Query Level ELO-DCG Algorithm

Input: Labeled set \mathcal{L} , unlabeled set \mathcal{U} , the number of queries to be selected n_q .

Output: The selected query set \mathcal{Q} .

Method:

```

for  $i = 1, \dots, N$  do  $N =$  size of the ensemble
  Subsample  $L$  and learn a relevance function
   $s_j^i \leftarrow$  score predicted by that function on the  $j$ th document
  in  $U$ .
end for
for  $q = 1, \dots, Q$  do  $Q =$  number of queries in  $U$ 
   $\mathcal{I} \leftarrow$  documents associated to  $q$ 
  for  $i = 1, \dots, N$  do
     $d_i \leftarrow$  BDCG( $\{G(s_j^i)\}_{j \in \mathcal{I}}$ )
  end for
   $t_j \leftarrow \langle G(s_j^i) \rangle$ 
   $d \leftarrow$  BDCG( $\{t_j\}_{j \in \mathcal{I}}$ )
   $EL(q) \leftarrow \langle d_i \rangle - d$ 
end for
Return  $\mathcal{Q}$  contains  $n_q$  queries with the highest values  $EL(q)$ .

```

4.2 Document Level Active Learning

Substituting (6) into (5), the expected loss of the i th document is expressed as:

$$\begin{aligned}
 EL(q, i) &= \int_{Y^i} \min_{\pi} \int_{y_i} (\max_{\pi'} DCG(\pi', Y) - DCG(\pi, Y)) dy_i dY^i, \\
 &= \int_{Y^i} \left[\int_{y_i} \max_{\pi'} DCG(\pi', Y) P(Y|X_q, D) dy_i \right. \\
 &\quad \left. - \max_{\pi} \int_{y_i} DCG(\pi, Y) P(Y|X_q, D) dy_i \right] dY^i,
 \end{aligned} \tag{10}$$

which is similar to equation (8) except that the uncertainty is on y_i instead of the entire vector Y and that there is an outer expectation on the relevance values for the other documents. The corresponding pseudo-code is provided in Algorithm 2.

Algorithm 2. Document Level ELO-DCG Algorithm

Input: Labeled set \mathcal{L} , unlabeled set \mathcal{U} , the number of documents to be selected n_u .

Output: The selected document set \mathcal{D} .

Method:

```

for  $i = 1, \dots, N$  do  $N =$  size of the ensemble
  Subsample  $L$  and learn a relevance function
   $s_j^i \leftarrow$  score predicted by that function on the  $j$ th document
  in  $U$ .
end for
for all  $j \in U$  do
  for  $i = 1, \dots, N$  do
     $t_k \leftarrow s_k^i, \forall k \neq j$ 
    for  $p = 1, \dots, N$  do
       $t_j \leftarrow s_j^p$ 
       $d_p \leftarrow$  BDCG( $\{G(t_k)\}$ )
    end for
     $g_k \leftarrow G(s_k^i), \forall k \neq j$ 
     $g_j \leftarrow \langle G(s_j^i) \rangle$ 
     $EL(j) \leftarrow \langle d_p \rangle - \text{BDCG}(\{g_k\})$ 
  end for
end for
Return  $\mathcal{D}$  that contains  $n_u$  documents with the highest values of  $EL(j)$ .

```

4.3 Two-Stage Active Learning

Both query level and document level active learning have their own drawbacks. Since query level active learning selects all documents associated with a query, it is tend to include non-informative documents when there are a large number of documents associated with each query. For example, in Web search applications, there are large amount of Web documents associated for a query; most of them are non-informative, since the quality of a ranking function is mainly measured by its ranking output on a small number of top ranked Web documents. On the other hand, document level active learning selects documents individually. This selection process implies unrealistic assumption that documents are independent, which leads to some undesirable results. For example, an informative query could be missed if none of its documents is selected; or only one document is selected for a query, which is not a good example in ranking learning.

Algorithm 3. Two Stage ELO-DCG Algorithm

Input: Labeled set \mathcal{L} , unlabeled doc \mathcal{U} , the number of queries to be selected n_q , the number of documents to be selected for each query n_u .

Output: The selected query-document set \mathcal{D} .

Method:

```

for  $i = 1, \dots, N$  do  $N =$  size of the ensemble
  Subsample  $L$  and learn a relevance function
   $s_j^i \leftarrow$  score predicted by that function on the  $j$ th document
  in  $U$ .
end for
for  $q = 1, \dots, Q$  do  $Q =$  number of queries in  $U$ 
   $\mathcal{I} \leftarrow$  documents associated to  $q$ 
  for  $i = 1, \dots, N$  do
     $d_i \leftarrow$  BDCG( $\{G(s_j^i)\}_{j \in \mathcal{I}}$ )
  end for
   $t_j \leftarrow \langle G(s_j^i) \rangle$ 
   $d \leftarrow$  BDCG( $\{t_j\}_{j \in \mathcal{I}}$ )
   $EL(q) \leftarrow \langle d_i \rangle - d$ 
end for
Select  $n_q$  queries with the highest values  $EL(q)$  into  $\mathcal{Q}$ .
for all  $q \in \mathcal{Q}$  do
   $\mathcal{I} \leftarrow$  documents associated to  $q$ 
  for all  $j \in \mathcal{I}$  do
    for  $i = 1, \dots, N$  do
       $t_k \leftarrow s_k^i, \forall k \neq j$ 
      for  $p = 1, \dots, N$  do
         $t_j \leftarrow s_j^p$ 
         $d_p \leftarrow$  BDCG( $\{G(t_k)\}$ )
      end for
       $g_k \leftarrow G(s_k^i), \forall k \neq j$ 
       $g_j \leftarrow \langle G(s_j^i) \rangle$ 
       $EL(j) \leftarrow \langle d_p \rangle - \text{BDCG}(\{g_k\})$ 
    end for
    Select  $n_u$  documents with the highest values of  $EL(j)$  into  $\mathcal{D}_q$ 
  end for
   $\mathcal{D} = \mathcal{D} \cup \mathcal{D}_q$ 
end for
Return  $\mathcal{D}$ .

```

Therefore, it is natural to combine query level and document level into two-stage active learning. A realistic




Fig. 1. Grade distribution for a randomly selected web search data set.

assumption for ranking data is that queries are independent and the documents are independent given on a query. Based on this assumption, we propose the two-stage ELO-DCG algorithm that is summarized in Algorithm 3. Algorithm 3 first selects most informative queries; then, selects the most informative documents for each selected query.

4.4 Balanced ELO Active Learning

In this section, we show how ELO framework can be easily adapted to address the grade imbalance problem for ranking.

In many ranking applications, especially Web search ranking, the distribution of relevance scores are very skewed. Under five grade scheme, "Perfect-4", "Excellent-3", "Good-2", "Fair-1", and "Bad-0", there are usually much less perfect and excellent examples in a randomly selected Web search data set from a commercial search engine. Fig. 1 shows the grade distribution of a randomly selected Web search data set. In Fig. 1, we observe that the excellent examples are less than 10 percent and perfect examples are less than 2 percent.

On the other hand, users usually only care about the top ranked documents, which are often perfect or excellent documents, instead of entire collection of documents that match a query. Intuitively, if a training set has very few perfect or excellent examples, it will be very difficult for the ranking learner to learn decision boundaries to identify perfect and excellent documents. As a result, the learned ranking function cannot be expected to have good performance. In fact, [2] shows that selecting more relevant documents into a training set can improve ranking models.

ELO framework can be easily to extended to address the skew grade distribution problem by re-designing the expected loss function. For the ELO-DCG algorithms, we propose the following expected loss function for the

document selection, which selects the examples with high expected DCG loss and at the same time put more preference on highly relevant documents to obtain more balanced train data,

$$\text{EL}(q, i) = \int_{Y^i} \min_{\pi} \int_{y_i} s(y_i) \ell(\pi, Y) P(Y|X_q, D) dy_i dY^i, \quad (11)$$

where $s(y_i)$ is a function to incorporate the prior knowledge to balance the grade distribution. For example,

$$s(y) = y \quad (12)$$

puts preference on highly relevant documents; and similarly,

$$s(y) = y^2 \quad (13)$$

put even more on highly relevant documents; and

$$s(y) = (y - \mathbb{E}(y))^2 \quad (14)$$

where \mathbb{E} denotes expectation, puts preference on both highly relevant documents and highly irrelevant documents.

In this paper we simply adopt the identity function in Eq.(12) to show how to deal with grade skew problem as shown in Fig. 1 under ELO framework. The approach can be easily extended to other situations in the need of balanced grade distribution.

Substituting (6) and (12) into (11), the expected loss of the i th document is expressed as:

$$\begin{aligned} \text{EL}(q, i) &= \int_{Y^i} \min_{\pi} \int_{y_i} y_i (\max_{\pi'} \text{DCG}(\pi', Y) - \text{DCG}(\pi, Y)) \\ &\quad dy_i dY^i, \\ &= \int_{Y^i} \left[\int_{y_i} y_i \max_{\pi'} \text{DCG}(\pi', Y) P(Y|X_q, D) dy_i \right. \\ &\quad \left. - \max_{\pi} \int_{y_i} y_i \text{DCG}(\pi, Y) P(Y|X_q, D) dy_i \right] dY^i, \end{aligned} \quad (15)$$

which is similar to equation (10) except that y_i is used to weight the loss function to put more weights on highly relevant documents.

When using (15) to select documents, we obtain balanced ELO-DCG algorithms. The balanced two stage ELO-DCG algorithm is summarized in Algorithm 4. Algorithm 4 first selects most informative queries as in Algorithm 2; then, selects the documents for each query by considering both DCG loss and grade distribution, i.e., selects documents with most potential to improve DCG and at the same time puts more preference on the highly relevant document to mitigate the skew grade distribution problem.

5 EXPERIMENTAL EVALUATION

As a general active learning algorithm for ranking, ELO-DCG can be applied to a wide range of ranking applications. In this section, we apply different versions of ELO-DCG algorithms to web search ranking to demonstrate the properties and effectiveness of our algorithm. We denote query level, document level, and two-stage ELO-DCG algorithms as ELO-DCG-Q, ELO-DCG-D, and ELO-DCG-QD, respectively.

Algorithm 4. Balanced two Stage ELO-DCG Algorithm

Input: Labeled set \mathcal{L} , unlabeled doc \mathcal{U} , the number of queries to be selected n_q , the number of documents to be selected for each query n_u .

Output: The selected query-document set \mathcal{D} .

Method:

for $i = 1, \dots, N$ **do** $N =$ size of the ensemble
 Subsample L and learn a relevance function
 $s_j^i \leftarrow$ score predicted by that function on the j th document in U .

end for

for $q = 1, \dots, Q$ **do** $Q =$ number of queries in U

$\mathcal{I} \leftarrow$ documents associated to q

for $i = 1, \dots, N$ **do**

$d_i \leftarrow$ BDCG($\{G(s_j^i)\}_{j \in \mathcal{I}}$)

end for

$t_j \leftarrow \langle G(s_j^i) \rangle$

$d \leftarrow$ BDCG($\{t_j\}_{j \in \mathcal{I}}$)

$EL(q) \leftarrow \langle d_i \rangle - d$

end for

Select n_q queries with the highest values $EL(q)$ into \mathcal{Q} .

for all $q \in \mathcal{Q}$ **do**

$\mathcal{I} \leftarrow$ documents associated to q

for all $j \in \mathcal{I}$ **do**

for $i = 1, \dots, N$ **do**

$t_k \leftarrow s_k^i, \forall k \neq j$

for $p = 1, \dots, N$ **do**

$t_j \leftarrow s_j^p$

$d_p \leftarrow$ BDCG($\{G(t_k)\}$)

end for

$g_k \leftarrow G(s_k^i), \forall k \neq j$

$g_j \leftarrow \langle G(s_j^i) \rangle$

$\bar{s}_j \leftarrow \langle s_j^p \rangle$

$EL(j) \leftarrow \bar{s}_j(\langle d_p \rangle - \text{BDCG}(\{g_k\}))$

end for

Select n_u documents with the highest values of $EL(j)$ into

\mathcal{D}_q

end for

$\mathcal{D} = \mathcal{D} \cup \mathcal{D}_q$

end for

Return \mathcal{D} .

5.1 Data Sets

We use web search data from a commercial search engine. The data set consists of a random sample of about 10,000 queries with about half million web documents. Those query-document pairs are labeled using a five-grade labeling scheme: {Bad, Fair, Good, Excellent, Perfect}.

For a query-document pair $(q; d)$, a feature vector x is generated and the features generally fall into the following three categories. Query features comprise features dependent on the query q only and have constant values across all the documents, for example, the number of terms in the query, whether or not the query is a person name, etc. Document features comprise features dependent on the document d only and have constant values across all the queries, for example, the number of inbound links pointing to the document, the amount of anchor-texts in bytes for the document, and the language identity of the document, etc. Query-document features comprise features dependent on the relation of the query q with respect to the document d ,

TABLE 1
Sizes of the Five Data Sets

Data set	Number of examples
base set 2k	~2,000
base set 4k	~4,000
base set 8k	~8,000
AL set	~160,000
test set	~180,000

for example, the number of times each term in the query q appears in the document d , the number of times each term in the query q appears in the anchor-texts of the document d , etc. We selected about five hundred features in total.

We randomly divide this data set into three subsets, base training set, active learning set, and test set. From the base training set, we randomly sample three small data sets to simulate small size labeled data sets \mathcal{L} . The active learning data set is used as a large size unlabeled data set \mathcal{U} from which active learning algorithms will select the most informative examples. The true labels from active learning set are not revealed to the ranking learners unless the examples are selected for active learning. The test set is used to evaluate the ranking functions trained with the selected examples plus base set examples. We kept test set large to have rigorous evaluations on ranking functions. The sizes of those five data sets are summarized in Table 1.

5.2 Experimental Setting

For the learner, we use gradient boosting decision tree (GBDT) [15].

The input for ELO-DCG algorithm is a base data set \mathcal{L} and the AL data set \mathcal{U} . The size of the function ensemble is set as eight for all experiments. ELO-DCG algorithm selects top m informative examples; those m examples are then added to the base set to train a new ranking function; the performance of this new function is then evaluated on the test set. Each algorithm with each base set is tested on 14 different m , 500, 1,000, 2,000, 4,000, 6,000, 8,000, 10,000, 12,000, 16,000, 24,000, 32,000, 48,000, 64,000, and 80,000. For every experimental setting, 10 runs are repeated and in each run the base set is re-sampled to generate a new function ensemble.

For the performance measure for ranking models, we select to use DCG-k, since users of a search engine are only interested in the top-k results of a query rather than a sorted order of the entire document collection. In this study, we select k as 10. The average DCG of 10 runs is reported for each experiment setting.

5.3 Document Level Active Learning

We first investigate document level active learning, since documents correspond to basic elements to be selected in the traditional active learning framework. We compare document level ELO-DCG algorithm with random selection (denoted by Random-D) and a classical active learning approach based on variance reduction (VR) [7], which selects document examples with largest variance on the prediction scores. Concretely, VR based approach first random sample data to train multiple models; then, it applies




Fig. 2. DCG comparison of document level ELO-DCG, variance reduction based document selection, and random document selection with base sets of sizes 2, 4, and 8k shows that ELO-DCG algorithm outperforms the other two document selection methods at various sizes of selected examples.

multiple models to the candidate document examples to obtain multiple prediction scores; finally, it selects the documents with largest variance on the prediction scores.

Fig. 2 compares the three document level active learning methods in terms of DCG-10 of the resulting ranking functions on the test set. Those ranking functions are trained with base data set and the selected examples. X-axis denotes number of examples selected by the active learning algorithm. For all three methods, the DCG increases with the number of added examples. This agrees with the intuition that the quality of a ranking function is positively correlated with the number of examples in the training set. ELO-DCG consistently outperforms the other two methods. A possible explanation is that ELO-DCG optimizes the expected DCG loss that is directly related to the objective function DCG-10 used to evaluate the ranking quality; on the other hand, the prediction score variance from VR is not directly related to DCG driven objective. In fact, VR performs even worse than random document selection when the size of the selected example is small. An advantage of the ELO-DCG algorithm is its capability to optimize directly based on the ultimate loss function to measure ranking quality.

5.4 Query Level Active Learning

In this section, we show that query level ELO-DCG algorithm effectively selects informative queries to improve the learning to rank performance. Since traditional active learning approaches cannot directly applied to query selection in ranking, we compare it with random query selection (denoted by Random-Q) used in practice.

Fig. 3 shows the DCG comparison results. we observe that for all three base sets, ELO-DCG performs better than random selection for all different sample sizes (from 500 to 80,000) that are added to base sets. Moreover, ELO-DCG converges much faster than random selection, i.e., ELO-DCG attains the best DCG that the whole AL data set can attain with much less examples added to the train data.

5.5 Two-Stage Active Learning

In this section, we compare two-stage ELO-DCG algorithm with other two two-stage active learning algorithms. One is two-stage random selection, i.e. random query selection followed by random document selection for each query. The other one is a widely used approach in practice, which first

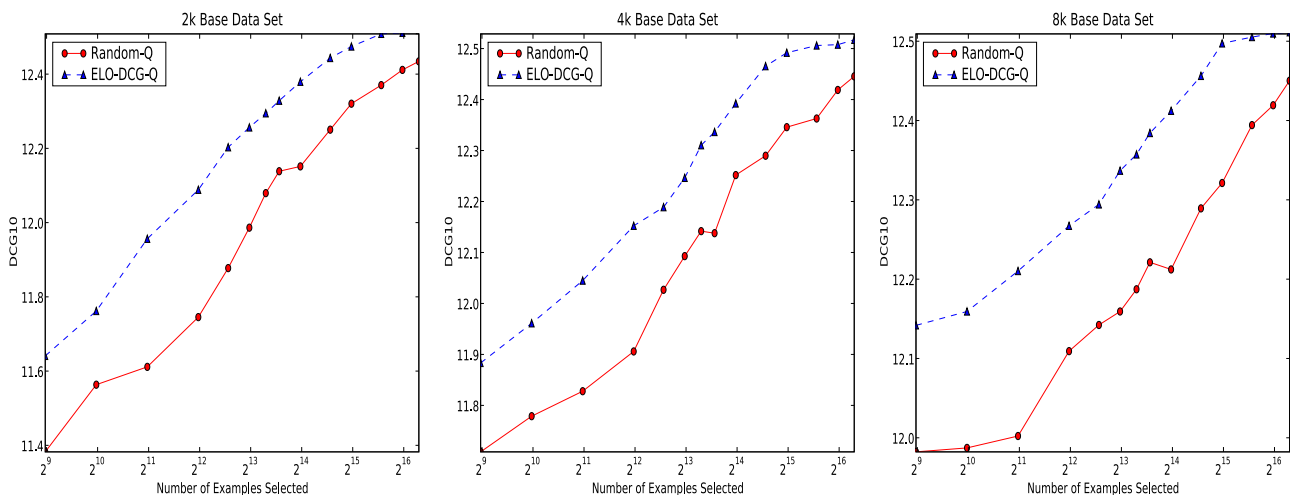


Fig. 3. DCG comparisons of query level ELO-DCG and random query selection with base sets of sizes 2, 4, and 8k shows that ELO-DCG algorithm outperforms random selection at various sizes of selected examples.

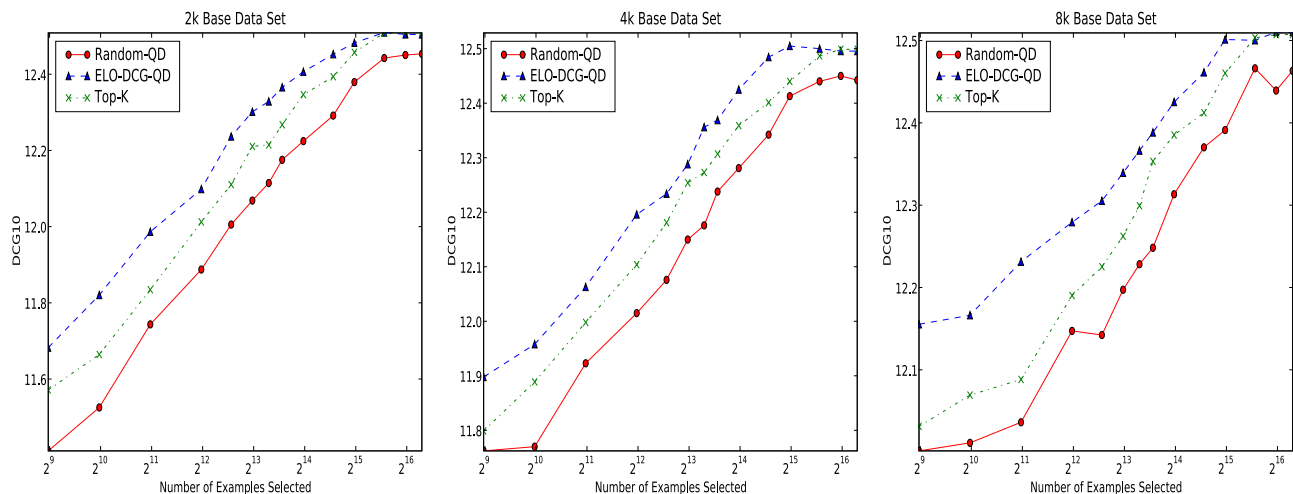


Fig. 4. DCG comparisons of two-stage ELO-DCG, two-stage random selection, and top-K selection with base set 2, 4, and 8k shows that ELO-DCG algorithm performs best.

randomly selects queries and then select top k relevant documents for each query based on current ranking functions (such as top k web sites returned by the current search engine) [24]. In our experimental setting, this approach corresponds to randomly query selection followed by selecting k documents with highest mean relevance scores within each selected query. We denote this approach as top-K. In all three two-stage algorithms, we simply fix the number of documents per query at 15 based on the results from [24].

Fig. 4 shows the DCG comparison results for two-stage active learning. We observe that among all three base sets, ELO-DCG performs the best and top-K performs the second. This result demonstrates that two-stage OLE-DCG effectively select most informative documents for most informative queries. A possible reason that top-K performs better than random selection is that top- k selects more perfect and excellent examples. Those examples contribute more to DCG than bad and fair examples.

We have observed that ELO-DCG algorithms perform best in all three active learning scenarios, query level, document level, and two stage active learning. Next, we compare three versions of ELO-DCG with each other.

Fig. 5 shows DCG comparisons of two-stage ELO-DCG, query level ELO-DCG, and document level ELO-DCG. We observe that for all three based sets, two stage ELO-DCG performs best. The reason that two-stage algorithm performs best may root in its reasonable assumption for the ranking data: queries are independent; the documents are conditionally independent given a query. On the other hand, the document level algorithm makes the incorrect assumption about document independence and may miss informative information at the query level; the query level algorithm selects all documents associated with a query, which are not all informative.

From the above results, we can observe that compared with the widely used top-K approach in practice, two-stage ELO-DCG algorithm can significantly reduce the labeling cost. Table 2 summarizes the cost reduction of two-stage algorithm compared with top-K approach.

In Table 2, the saturated size means that when the examples of this size are selected and added back to the base set to train a ranking function, the performance of the learned ranking function is equivalent to the ranking function trained with all active learning data. From the first row of




Fig. 5. DCG comparisons of two-stage ELO-DCG, query level ELO-DCG, and document level ELO-DCG, with base sets of sizes 2, 4, and 8k shows that two-stage ELO-DCG algorithm performs best.

TABLE 2
Cost Reduction of Two-Stage Algorithm Compared with Top-K Approach

Base set size	Saturated size for ELO-DCG	Saturated size for top-K	Percentage of cost reduced
2k	64k	80k	20%
4k	48k	80k	40%
8k	32k	80k	64%




Fig. 6. DCG comparisons of two-stage ELO-DCG and balanced two-stage ELO-DCG, with base sets of sizes 2, 4, and 8k shows that balanced two-stage ELO-DCG algorithm performs better.

Table 2, we observe that for the base set of size 2k, 64k is the saturated size for two-stage ELO-DCG algorithm and 80k is the saturated size for top-K approach; hence, 64k selected examples from two-stage ELO-DCG algorithm is equivalent to 80k selected examples from top-K approach. This means that two-stage ELO-DCG algorithm can reduce the labeling cost by 16k, which is 20 percent of 80k. The largest percentage of cost reduced, 64, percent is from base set 8k.

5.6 Balanced Two-Stage Active Learning

In this section, we compare balanced two-stage ELO-DCG algorithm (denoted as B-ELO-DCG) with normal two-stage ELO-DCG algorithm.

Fig. 6 shows DCG comparisons of balanced two-stage ELO-DCG, and normal two stage ELO-DCG. We observe that for all three based sets, balanced two stage ELO-DCG performs better. Grade distribution analysis of the examples selected by two algorithms shows that averagely the examples selected by balanced ELO-DCG algorithm have 10.1 percent more excellent documents and 5.2 percent more perfect documents, i.e., the data sets from balanced ELO-DCG have more balanced grade distributions toward to highly relevant documents. Balanced ELO-DCG algorithm selects examples with both high potential to improve DCG and more balanced grade distribution, which helps learning algorithm to learn better decision boundaries for identify highly relevant documents; hence, it leads better DCG performance.

6 CONCLUSIONS AND FUTURE WORK

We propose a general expected loss optimization framework for ranking, which is applicable to active learning

scenarios for various ranking learners. Under ELO framework, we derive novel algorithms, query level ELO-DCG and document level ELO-DCG, to select most informative examples to minimize the expected DCG loss. We propose a two stage active learning algorithm to select the most effective examples for the most effective queries. We further extend the proposed algorithm to deal with the typical skew grade distribution problem in learning to rank. Extensive experiments on real-world web search data sets have demonstrated great potential and effectiveness of the proposed framework and algorithms.

In future, we will investigate how to fuse the query level and document level selection steps in order to produce a more robust query selection strategy. Besides, we will also evaluate our active learning method upon different types of data.

ACKNOWLEDGMENTS

This research was partially supported by National Natural Science Foundation of China (No. 61003107; No. 61221001) and the High Technology Research and Development Program of China (No. 2012AA011702).


REFERENCES

- [1] N. Abe and H. Mamitsuka, "Query learning strategies using boosting and bagging," in *Proc. 15th Int. Conf. Mach. Learn.*, 1998, pp. 1–9.
- [2] J. A. Aslam, E. Kanoulas, V. Pavlu, S. Savev, and E. Yilmaz, "Document selection methodologies for efficient and effective learning-to-rank," in *Proc. 32nd Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2009, pp. 468–475.
- [3] J. Berger, *Statistical Decision Theory and Bayesian Analysis*. New York, NY, USA: Springer, 1985.
- [4] C. Campbell, N. Cristianini, and A. Smola, "Query learning with large margin classifiers," in *Proc. 17th Int. Conf. Mach. Learn.*, 2000, pp. 111–118.

- [5] B. Carterette, J. Allan, and R. Sitaraman, "Minimal test collections for retrieval evaluation," in *Proc. 29th Annu. Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2006, pp. 268–275.
- [6] W. Chu and Z. Ghahramani, "Extensions of Gaussian processes for ranking: Semi-supervised and active learning," in *Proc. Nips Workshop Learn. Rank*, 2005, pp. 33–38.
- [7] D. A. Cohn, Z. Ghahramani, and M. I. Jordan, "Active learning with statistical models," in *Proc. Adv. Neural Inf. Process. Syst.*, 1995, vol. 7, pp. 705–712.
- [8] D. Cossack and T. Zhang, "Subset ranking using regression," in *Proc. 19th Annu. Conf. Learn. Theory*, 2006, pp. 605–619.
- [9] I. Dagan and S. P. Engelsson, "Committee-based sampling for training probabilistic classifiers," in *Proc. 12th Int. Conf. Mach. Learn.*, 1995, pp. 150–157.
- [10] P. Donmez and J. Carbonell, "Active sampling for rank learning via optimizing the area under the ROC curve," in *Proc. 31th Eur. Conf. IR Res. Adv. Inform. Retrieval*, 2009, pp. 78–89.
- [11] P. Donmez and J. G. Carbonell, "Optimizing estimated loss reduction for active sampling in rank learning," in *Proc. 25th Int. Conf. Mach. Learn.*, 2008, pp. 248–255.
- [12] V. Fedorov, *Theory of Optimal Experiments*. New York, NY, USA: Academic, 1972.
- [13] Y. Freund, R. Iyer, R. E. Schapire, and Y. Singer, "An efficient boosting algorithm for combining preferences," *J. Mach. Learn. Res.*, vol. 4, pp. 933–969, 2003.
- [14] Y. Freund, H. S. Seung, E. Shamir, and N. Tishby, "Selective sampling using the query by committee algorithm," *Mach. Learn.*, vol. 28, nos. 2–3, pp. 133–168, 1997.
- [15] J. Friedman, "Greedy function approximation: A gradient boosting machine," *Ann. Statist.*, vol. 29, pp. 1189–1232, 2001.
- [16] T. Fushiki, "Bootstrap prediction and Bayesian prediction under misspecified models," *Bernoulli*, vol. 11, no. 4, pp. 747–758, 2005.
- [17] D. D. Lewis and W. A. Gale, "A sequential algorithm for training text classifiers," in *Proc. 17th Annual Int. ACM SIGIR Conf. Research Development Info. Retrieval (SIGIR '9)*, 1994, pp. 3–12.
- [18] D. Lewis and W. Gale, "Training text classifiers by uncertainty sampling," in *Proc. 17th Annu. Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 1994, pp. 3–12.
- [19] B. Long, O. Chapelle, Y. Zhang, Y. Chang, Z. Zheng, and B. Tseng, "Active Learning for Ranking through Expected Loss Optimization," in *Proc. 33rd Annu. ACM SIGIR Conf.*, 2010, pp. 267–274.
- [20] A. McCallum and K. Nigam, "Employing EM and pool-based active learning for text classification," in *Proc. 5th Int. Conf. Mach. Learn.*, 1998, pp. 359–367.
- [21] B. Settles, "Active learning literature survey," *Comput. Sci. Dept., Univ. of Wisconsin, Madison, WI, USA, Tech. Rep. 1648*, 2009.
- [22] F. Xia, T.-Y. Liu, J. Wang, W. Zhang, and H. Li, "Listwise approach to learning to rank: theory and algorithm," in *Proc. 25th Int. Conf. Mach. Learn.*, 2008, pp. 1192–1199.
- [23] L. Yang, L. Wang, B. Geng, and X.-S. Hua, "Query sampling for ranking learning in web search," in *Proc. 32nd Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2009, pp. 754–755.
- [24] E. Yilmaz and S. Robertson, "Deep versus shallow judgments in learning to rank," in *Proc. 32nd Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2009, pp. 662–663.
- [25] H. Yu, "SVM selective sampling for ranking with application to data retrieval," in *Proc. 11th ACM SIGKDD Int. Conf. Knowl. Discovery. Data Mining*, 2005, pp. 354–363.
- [26] Z. Zheng, H. Zha, T. Zhang, O. Chapelle, K. Chen, and G. Sun, "A general boosting method and its application to learning ranking functions for web search," in *Proc. Adv. Neural Inf. Process. Syst.*, 2008, pp. 1697–1704.
- [27] O. Zoeter, N. Craswell, M. Taylor, J. Guiver, and E. Snelson, "A decision theoretic framework for implicit relevance feedback," in *Proc. NIPS Workshop Mach. Learn. Web Search*, 2007, pp. 133–142.
- [28] T. Joachims, "Optimizing search engines using clickthrough data," in *Proc. 8th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, 2002, pp. 133–142.
- [29] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender, "Learning to rank using gradient descent," in *Proc. 22nd Int. Conf. Mach. Learn.*, 2005, pp. 89–96.
- [30] Y. Freund, R. Iyer, R. Schapire, and Y. Singer, "An efficient boosting algorithm for combining preferences," in *Proc. 15th Int. Conf. Mach. Learn.*, 1998, pp. 391–398.
- [31] J. Xu and H. Li, "Adarank: A boosting algorithm for information retrieval," in *Proc. 30th Annu. Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2007, pp. 391–398.
- [32] C. Cortes, M. Mohri, and A. Rastogi, "Magnitude-preserving ranking algorithms," in *Proc. 24th Int. Conf. Mach. Learn.*, 2007, pp. 169–176.
- [33] Z. Zheng, K. Chen, G. Sun, and H. Zha, "A regression framework for learning ranking functions using relative relevance judgments," in *Proc. 30th Annu. Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2007, pp. 287–294.
- [34] J. Guiver and E. Snelson, "Learning to rank with SoftRank and Gaussian processes," in *Proc. 31st Annu. Int. ACM SIGIR Conf. Res. Develop. Inform. Retrieval*, 2008, pp. 259–266.
- [35] S. Rodrigo, M. A. Goncalves, and A. Veloso, "Rule-based active sampling for learning to rank," in *Proc. Eur. Conf. Mach. Learn. Knowl. Discovery Databases*, 2011, pp. 240–255.
- [36] B. Qian, H. Li, J. Wang, X. Wang, and I. Davidson, "Active Learning to Rank using Pairwise Supervision," in *Proc. 13th SIAM Int. Conf. Data Mining*, 2013, pp. 297–305.



Bo Long is a staff applied researcher at LinkedIn Inc, and was formerly a senior research scientist at Yahoo! Labs. His research interests include data mining and machine learning with applications to web search, recommendation, and social network analysis. He holds eight innovations and has published peer-reviewed papers in top conferences and journals including ICML, KDD, ICDM, AAAI, SDM, CIKM, and KAIS. He has served as a reviewer, workshops co-organizer, conference organizer committee member, and area chair for multiple conferences, including KDD, NIPS, SIGIR, ICML, SDM, CIKM, JSM, etc.




Jiang Bian received the PhD degree in computer science from the Georgia Institute of Technology, Atlanta GA, in 2010. He is a researcher at Microsoft Research. Prior to that, he was a scientist at Yahoo! Labs. His research interests include machine learning, information retrieval, data mining, social network analysis, and computational advertising. More details of his research and background can be found at <https://sites.google.com/site/jiangbianhome/>.



Olivier Chapelle received his graduation in theoretical computer science from the Ecole Normale Supérieure de Lyon in 1999 and the PhD degree from the University of Paris in 2002. He is a principal research scientist at Criteo, where he works on machine learning for display advertising. Prior to that, he was part of the machine learning group of Yahoo! Research and before that he worked at the Max Planck Institute in Tbingen. His main research interests include kernel machines, semi-supervised learning, ranking and large scale learning. He has published more than 80 publications and has been granted more than 10 patents. He has served as an associated editor for the *Machine Learning Journal and Transactions on Pattern Analysis and Machine Intelligence*.



Ya Zhang received the BS degree from Tsinghua University and the PhD degree in information sciences and technology from the Pennsylvania State University. Since March 2010, she has been a professor in the Institute of Image Communication and Networking at Shanghai Jiao Tong University. Prior to that, she worked at Lawrence Berkeley National Laboratory, University of Kansas, and Yahoo! Labs. Her research interest include data mining and machine learning, with applications to information retrieval, web mining, and multimedia analysis.



Yoshiyuku Inagaki received the PhD degree in social sciences, specialized on mathematical behavioral science. He is a research scientist at Yahoo! Labs. Most of his research centers around machine learned ranking. Before joining Yahoo! Labs, he worked on medical image analysis systems at the University of California, Irvine. His current research interests are machine learning, deep learning, topic modeling, and artificial intelligence.



Yi Chang is a director of science in Yahoo Labs, where he is managing the search and anti-abuse science team. His research interests include information retrieval, applied machine learning, natural language processing, and social media mining. He has published more than 70 referred journal and conference publications, and he is a co-author of *Relevance Ranking for Vertical Search Engines*, which is the first book of this area.

▷ **For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.**