

SetRank: Learning a Permutation-Invariant Ranking Model for Information Retrieval

Liang Pang¹, Jun Xu^{2,3,*}, Qingyao Ai⁴, Yanyan Lan¹, Xueqi Cheng¹, Jirong Wen^{2,3}

¹CAS Key Lab of Network Data Science and Technology,

Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

²Gaoling School of Artificial Intelligence, Renmin University of China

³Beijing Key Laboratory of Big Data Management and Analysis Methods ⁴The University of Utah, USA

{pangliang,lanyanyan,cxq}@ict.ac.cn,{junxu,jrwen}@ruc.edu.cn,aiqy@cs.utah.edu

ABSTRACT

In learning-to-rank for information retrieval, a ranking model is automatically learned from the data and then utilized to rank the sets of retrieved documents. Therefore, an ideal ranking model would be a mapping from a document set to a permutation on the set, and should satisfy two critical requirements: (1) it should have the ability to model cross-document interactions so as to capture local context information in a query; (2) it should be permutation-invariant, which means that any permutation of the inputted documents would not change the output ranking. Previous studies on learning-to-rank either design uni-variate scoring functions that score each document separately, and thus failed to model the cross-document interactions; or construct multivariate scoring functions that score documents sequentially, which inevitably sacrifice the permutation invariance requirement. In this paper, we propose a neural learning-to-rank model called SetRank which directly learns a *permutation-invariant ranking model defined on document sets of any size*. SetRank employs a stack of (induced) multi-head self attention blocks as its key component for learning the embeddings for all of the retrieved documents jointly. The self-attention mechanism not only helps SetRank to capture the local context information from cross-document interactions, but also to learn permutation-equivariant representations for the inputted documents, which therefore achieving a permutation-invariant ranking model. Experimental results on three benchmarks showed that the SetRank significantly outperformed the baselines include the traditional learning-to-rank models and state-of-the-art Neural IR models.

CCS CONCEPTS

• Information systems → Learning to rank.

KEYWORDS

Learning to rank; permutation-invariant ranking model

*Corresponding author

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '20, July 25–30, 2020, Virtual Event, China

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-8016-4/20/07...\$15.00

<https://doi.org/10.1145/3397271.3401104>

ACM Reference Format:

Liang Pang, Jun Xu, Qingyao Ai, Yanyan Lan, Xueqi Cheng, Jirong Wen. 2020. SetRank: Learning a Permutation-Invariant Ranking Model for Information Retrieval. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR'20)*, July 25–30, 2020, Virtual Event, China. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3397271.3401104>

1 INTRODUCTION

Learning-to-rank has been extensively studied in both academia and industry. Usually, the task of learning-to-rank can be described as two steps. First, the training step, where a ranking model that projects each query-document pair to a ranking score is constructed and learned from labeled data such as user clicks and relevance annotations. And second, the testing step, where the learned ranking model is applied to a set of documents retrieved for a new query and finally returns a ranked document list to the users.

Traditional learning-to-rank models are usually designed on the basis of the probability ranking principle (PRP) [33]. PRP assumes that each document has a unique probability to satisfy a particular information need. Therefore, the ranking scores of documents are assigned separately and are independent to each other. Despite widely adopted, the power of PRP-based learning-to-rank methods, however, have been proved to be limited [1, 2, 22]. First, independent scoring paradigms prevent traditional learning-to-rank models from modeling cross-document interactions and capturing local context information. As shown by previous studies on pseudo relevance feedback [26] and query-dependent learning-to-rank [9], incorporating local context information such as query-level document feature distributions can significantly improve the effectiveness of modern ranking systems. Second, as pointed by Robertson [33], PRP works document-by-document while the results of ranking should be evaluated request-by-request. Behavior analysis on search engine users manifest that user's interactions with information retrieval systems show strong *comparison* patterns [24, 40]. In practice, search engine users often compare multiple documents on a result page before generating a click action. Also, studies on query-document relevance annotations show that information from other documents in the same ranked list could affect an annotator's decision on the current document [34, 39], which challenge the basic hypothesis that relevance should be modeled independently on each document for a single information request.

Recently, a new group of learning-to-rank methods whose scoring functions take multiple documents as input and jointly predict their ranking scores have received more and more attention. For

example, Ai et al. [1, 2] propose to learn context-aware query-level ranking model that takes a sequence of documents as its input. By modeling and comparing multiple documents together, a new scoring functions, namely the *multivariate scoring functions*, naturally capture the local context information and produce the state-of-the-art performance on many learning-to-rank benchmarks. Nonetheless, existing multivariate scoring approaches is sensitive to the order of the inputted documents, which violates the *permutation invariance* requirement for ranking model. That is because the models used by previous works, e.g. dense neural network [2] and recurrent neural network [1], all assume that the input is a sequence of documents and heavily biased to its initial ranking. When the initial ranking has poor performance or disturbed by accident, the performances of these models decline significantly.

Inspired by the work of Set Transformer [27], in this paper, we propose to develop a multivariate ranking model whose input is a document set of any size, and the output is a permutation over the set. The model, referred to as SetRank, is to learn a multivariate scoring function mapping from document set to its permutation, making full use of self-attention mechanism [35]. Firstly, It receives a whole set of documents (encoded as a set of vectors) as input. Then, employs a stack of (induced) multi-head self attention blocks as its key component for learning the embeddings for the whole set of documents jointly. Finally, the output of the last block is considered as the representation of the documents, and a row-wise fully connected network is used to generate the ranking scores.

SetRank offers several advantages. First of all, similar to existing work on multivariate scoring functions, SetRank considers the inputted documents as a whole via self-attention mechanism and model the interrelationship between them in ranking, which makes them more competitive than traditional learning-to-rank methods based on univariate scoring functions. Secondly, as shown in the theoretical analysis, SetRank learns a permutation-invariant function as its ranking model, that is, any permutation of the inputted documents does not change the outputted document ranking. Additionally, self-attention is an additive function which is insensitive to the size of document set, while a sequential model, such as recurrent neural network, using nested function is severely affected by the size of document set. Compared to existing learning-to-rank models which learn multivariate ranking functions defined on pre-defined document sequences, SetRank is more natural for document ranking and more friendly for parallel computing. Last but not least, with the help of ordinal embeddings, SetRank is enabled to involve multiple document rankings as the initial rankings.

Experimental results on three large scale publicly available benchmarks showed that SetRank significantly outperformed the base-lines including the traditional learning-to-rank models, such as RankSVM [23], RankBoost [15], Mart [16], and LambdaMART [8], and state-of-the-art neural IR models, such as DLCM [1] and GSF [2]. Analyses showed that the ranker learned by SetRank is robust to the input orders and sizes, making it to be a stable ranking model.

2 RELATED WORK

This section reviews previous studies on learning-to-rank, deep learning for ranking, and the machine learning defined on sets.

2.1 Learning-to-Rank

learning-to-rank refers to a group of algorithms that apply machine learning techniques to solve ranking problems. It usually represents each query-document pair with a feature vector created by human or extracted with some data pre-processing techniques. Then, the feature vectors are fed into a machine learning model to produce scores for each document so that the final ranking can be generated by sorting documents according to their scores.

In the past two decades, learning-to-rank methods based on different machine learning algorithms [7, 16, 23] has been proposed and applied to a variety of ranking tasks including document retrieval [28], question answering [38], recommendation [13], conversational search [43], etc. Based on their loss functions, existing learning-to-rank algorithms can be categorized into three groups: pointwise [16], pairwise [7, 23], and listwise [5, 6, 10, 37]. Despite their differences on loss functions, most traditional learning-to-rank methods are constructed based on the probability ranking principle [33], which assumes that documents should be ranked individually according to their probabilities of satisfying a particular information need. Under this assumption, they compute the ranking scores of each document separately based on their feature vectors, which is referred to as the uni-variate scoring paradigm. Through intuitive, the uni-variate scoring paradigm is sub-optimal because it limits the model's power on capturing inter-document relationships and local context information, which is important for understanding user relevance in online systems [2].

To solve this problem, previous studies propose a new ranking paradigm that directly rank a list of documents together with a multivariate scoring function. For example, Ai et al. [1] propose to encode the local context of an initial rank list with a recurrent neural network and rerank the documents based on the latent context embedding of all documents. Bello et al. [4], Jiang et al. [22] propose a slate optimization framework that directly predict the ranking of a list of documents by jointly considering their features together. Ai et al. [2] further formalize the idea and propose a multivariate scoring framework based on deep neural networks. [31] uses self-attention network to aggregate cross-documents information. In recommendation, Pei et al. [32] propose to construct a list-wise re-ranking framework that personalizes recommendation results by jointly considering multiple items in the same ranked list.

In this work, we propose a new learning-to-rank model based on the multivariate scoring paradigm. Different from previous studies that conduct reranking based on an initial order of documents [1, 2, 4], our method is permutation-invariant and can be applied directly to a set of documents with or without any preprocessing.

2.2 Deep Learning for Ranking

Recently, deep learning methods such as deep neural networks and convolution neural networks have been widely applied to IR problems. They have been proved to be effective in capturing latent semantics and extracting effective features for ranking problems. For example, Huang et al. [20] propose a deep structured semantic model for ad-hoc retrieval that directly predict the relevance of each document based on a multi-layer feed-forward network and the embedding of trigrams in document titles. Guo et al. [18] analyze the key factors of relevance modeling and propose a deep relevance

matching model that ranks documents based on the matching histogram of each query and document in the latent semantic space. Later, Pang et al. [30] propose a DeepRank model that jointly considers the strength and patterns of query-document matching and uses a recurrent neural network to measure and aggregate the local relevance information for ranking. Another type of deep models, originated from NLP, treat retrieval problems as a text matching between queries and documents. For example, Hu et al. [19] propose a group of ARC models that build text representations based on word embeddings and convolution neural networks for matching proposes. Pang et al. [29] propose to match text based on a word matching matrix and propose a MatchPyramid model for paraphrase identification and question answering. Wan et al. [36] further propose a Match-SRNN model that uses a spatial recurrent neural network to capture the recursive matching structure in text matching problems.

The proposed model in this paper can be considered as a deep model for IR because we take the self-attention networks [35], a popular neural technique used in machine learning tasks [12, 42], as our building blocks. Different from previous models that directly extract features from raw text, we focus on a general task of learning-to-rank and assume that feature representation for each query-document pair has been computed in advance. However, it's worth noting that our model framework also supports the joint learning of features and ranking models. Also, in spite of their structures, most existing deep learning models for IR is confined to the same uni-variate scoring paradigm used by traditional learning-to-rank algorithms, while our model is a listwise ranking model under the multivariate scoring paradigm.

2.3 Machine Learning Defined on Sets

Typical machine learning algorithms such as regression or classification are designed for processing fixed dimensional data instances. Recent studies showed that they can be extended to handle the case when the inputs are permutation-invariant sets [14, 27, 41]. For example, Edwards and Storkey [14] demonstrated neural statistician which is an extension of variational autoencoder and can learn statistics of datasets. Zaheer et al. [41] theoretically characterize the permutation-invariant functions and provide a family of functions to which any permutation-invariant objective function must belong. Lee et al. [27] further present the Set Transformer model to capture the interactions among elements in the input set. In this paper, we borrow the self-attention mechanism proposed in Transformer to capture the interactions among documents in ranking.

3 RANKING MODEL DEFINED ON SETS

In this section, we formalize document ranking as learning a multivariate scoring function defined on document sets, and then analyze the requirements of the function.

3.1 Problem Formulation

Generally speaking, the task of ranking aims to measure the relative order among a set of items under specific constraints. When ranking is completed, the arranged items is returned, called a permutation. When being applied to learning-to-rank for information retrieval,

the constraint corresponds to a query q , and the set of items correspond to a set of N documents $D = [d_i]_{i=1}^N \subseteq \mathbb{D}$ retrieved by the query q , and \mathbb{D} is the set of all indexed documents. The task of ranking is to find a permutation $\pi \in \Pi_N$ on the document set D so that some utility is maximized, where Π_N is the set of all permutations of indices $\{1, 2, \dots, N\}$.

In learning-to-rank, the ranking models are trained with a set of labeled query-document pairs. For each query q , its retrieved documents and the corresponding relevance labels are provided and denoted as $\psi_q = \{D = \{d_i\}, \mathbf{y} = \{y_i\} | 1 \leq i \leq N\}$, where y_i denotes the relevance label corresponds to d_i . Therefore, the dataset that contains all of the training queries can be denoted as $\Psi = \{\psi_q\}$. The goal of learning-to-rank is to minimize the empirical loss over the training data and to yield the optimal scoring function $F(\cdot)$:

$$\mathcal{L}(F) = \frac{1}{|\Psi|} \sum_{(D, \mathbf{y}) \in \Psi} l(\mathbf{y}, F(D)), \quad (1)$$

where $l(\cdot)$ is the loss function; $F(\cdot) : \mathbb{D}^N \mapsto \mathbb{R}^N$ is the scoring function which is responsible for assigning scores to the documents, so that the result permutation on D can be generated by sorting with the scores: $\hat{\pi} = \text{sort} \circ F(D)$. The operator 'sort' sorts the documents in D according to the scores assigned by $F(D)$ in descending order.

Most of the ranking models learnt by the existing learning-to-rank models (including point-wise, pair-wise, and list-wise models) are uni-variate scoring functions $g : \mathbb{D} \mapsto \mathbb{R}$. They are special cases of F because F can be written as

$$F(D)|_i = g(d_i), 1 \leq i \leq N, \quad (2)$$

where $F(D)|_i$ is the i -th dimension of $F(D)$. It is a natural and direct choice that fits the probability ranking principle [33].

3.2 Requirements of Ranking Model

We now show that a scoring function F defined on document sets should satisfy two critical requirements. First, it should be able to characterize the inter-relationship among the retrieved documents, called cross-document interaction. Second, the output of the function should not be changed under any permutation of the documents in the input set, called permutation invariance.

3.2.1 Cross-document Interactions. Although the PRP principle [33] and the derived uni-variate scoring functions have been widely used in information retrieval, recent studies show that they have several limitations in real settings. The major assumption in PRP is that the relevance of a document to a query is independent of the relevance of other documents. However, Fuhr [17] pointed out that "the relevance of any additional relevant document clearly depends on the relevant documents seen before"; Robertson [33] also pointed out that "PRP works document-by-document while the results of ranking should be evaluated request-by-request"; Scholer et al. [34], Yang [39] studied the relevance annotations and showed that information from other documents in the same ranked list could affect an annotator's decision on the current document.

Therefore, the assumptions underlying the classical PRP and the uni-variate scoring function shown in Equation (2), is hard to hold in some settings. It is expected that the multivariate scoring function F can deal with a more general ranking problem of taking account of the dependency information and cross-document interactions.

3.2.2 Permutation Invariance. The input to a ranking model is a set that contains multiple documents. According to the definition of a set, the response of a function defined on a set should be “indifferent” to the ordering of the elements. That is, the ranking model should satisfy the permutation invariance property: the target ranking for a given set is the same regardless of the order of documents in the set. That is, if we exchange the input positions of d_i and d_j , it will not affect the result of the final ranking. In practice, input sequence often convey covert bias information, e.g. order by document IDs, order by create time, or order by time-consuming, which may not correlate with the usefulness of the documents and could reduce the robustness of permutation awareness model. A permutation-invariant model addresses this problem with theoretical guarantees.

One way to construct a permutation-invariant ranking model is first assigning relevance scores to the documents with a permutation equivariant scoring function, and then sorting according to the scores. Lee et al. [27] give a definition of permutation equivariant scoring function, as shown in Definition 1.

DEFINITION 1. [27] Let Π_N be the set of all permutations of indices $\{1, \dots, N\}$. A function $F : X^N \rightarrow Y^N$ is permutation equivariant if and only if $\forall \pi \in \Pi_N$,

$$F([d_{\pi(1)}, \dots, d_{\pi(N)}]) = [F(D)|_{\pi(1)}, \dots, F(D)|_{\pi(N)}],$$

where $D = [d_1, \dots, d_N]$ is a set with the object order $1, 2, \dots, N$, and $F(D)|_{\pi(i)}$ is the $\pi(i)$ -th dimension of $F(D)$.

An example of a permutation equivariant function in IR ranking is the uni-variate scoring function shown in Equation (2).

4 OUR APPROACH: SETRANK

In this section, we penetrate into the details of our proposed SetRank model, and introduce the training method.

4.1 Overall Architecture

Inspired by the Transformer for machine translation [11, 35] and language modeling [12], we have devised a novel ranking model defined on document sets and satisfy the above requirements. The model is referred to as ‘SetRank’ and is shown in Figure 1. SetRank makes use of multi-head self attention blocks and its modifications to construct a multivariate scoring function that satisfies the *cross-document interactions* and the *permutation invariance* requirements, additionally adapts to the set input of any size (*set size adaptation*).

The pipeline of document ranking in SetRank consists of three layers: representation, encoding, and ranking. First, the representation layer separately represents each inputted document as a vector of features such as the hand-crafted features used in traditional learning-to-rank models. Also, it can involve the initial ranking of the documents in feature representations through the ordinal embeddings. The initial rankings could be generated by existing ranking models such as BM25 or a trained LambdaMART model. Second, the encoding layer enriches each query-document pair feature vector by involving other feature vectors of the associated documents. In this paper, we use Multi-head Self Attention Block (MSAB) or Induced Multi-head Self Attention Block (IMSAB) to take a set of query-document pairs representations as their input and generate a set of new representations using the self-attention blocks. Multiple sub-layers of MSAB or IMSAB blocks are stacked together

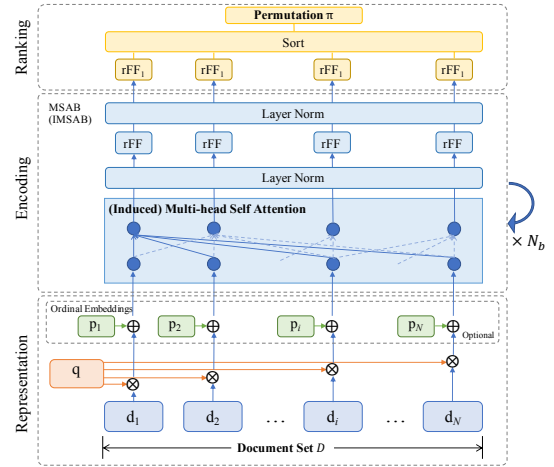


Figure 1: Architecture of SetRank. The representation layer generates representation of query-document pairs separately; the encoding layer jointly process the documents with multiple sub-layers of MSAB (or IMSAB), and output internal document representations; the ranking layer calculates the scores and sorts the documents.

with identical structure for modeling the high-order interactions between the documents. Third, the ranking layer receives the output vectors of the top-most MSAB (or IMSAB) block, passes them to a row-wise feed-forward (rFF) function, generates relevance scores for all of the documents, and finally sorts the documents according to these scores. The following sections will introduce the details of the three layers.

4.2 Document Representation

Given a query q and its associated document set $D = [d_1, d_2, \dots, d_N]$, each of the document in D can be represented as a feature vector

$$\mathbf{d}_i = \phi(q, d_i), \text{ where } \mathbf{d}_i \in \mathbb{R}^E,$$

where ϕ is the function for feature extraction and E is the dimension of the vector. The features extracted in traditional learning-to-rank are used here, including document only feature of PageRank, query-document matching features of TF-IDF, BM25 etc. In the experiments of this paper, we used the features provided by the benchmark datasets.

Besides the traditional learning-to-rank features, SetRank can optionally include the document ordinal embeddings as the input. In real search engines, the associated documents may have some prior rankings generated by the default ranking models such as BM25 or LambdaMART etc. To involve these initial ranking information and inspired by the positional embedding in Transformer [35], we propose an ordinal embedding function P which takes the absolute ranking position of a document as input, and encodes the position to a vector as the same dimension with \mathbf{d}_i :

$$\mathbf{p}_i = P(\text{rank}(d_i)), \text{ where } \mathbf{p}_i \in \mathbb{R}^E,$$

where $\text{rank}(d_i)$ denotes the absolute rank position of d_i in the initial ranking generated by models such as BM25, LambdaMART etc.

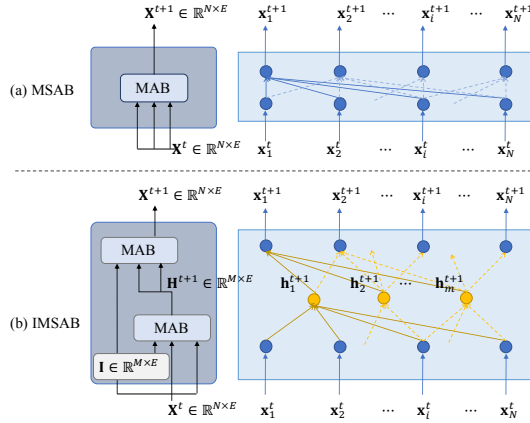


Figure 2: (a) The structure of multi-head self attention block (MSAB). MSAB encodes an arbitrary set X^t of size N and output an set X^{t+1} of size N . (b) The structure of induced multi-head self attention block (IMSAB) with induced size of M . IMSAB first encodes an arbitrary set X^t of size N to a fixed set H^{t+1} of size M , and then encodes H^{t+1} to an N -size output X^{t+1} . IMSAB can be interpreted as first creating M cluster centers and then encoding the inputted N documents using these M cluster centers.

The ordinal embedding vectors and the learning-to-rank feature vectors are respectively added, forming a feature matrix $X \in \mathbb{R}^{N \times E}$ for the N retrieved documents:

$$X = [d_1 + p_1, d_2 + p_2, \dots, d_N + p_N]^T.$$

Note that in some cases we may have more than one initial rankings because different ranking model (e.g., BM25 and LM4IR) may be applied simultaneously. The ordinal embeddings of a document corresponding to different rankings can be summed together, forming an overall ordinal embedding vector. In this way, SetRank is enabled to input multiple initial rankings. Also note that though the ordinal embedding mechanism is similar to the position embedding in Transformer, they are different in that the position embedding characterizes the position of a word in a sequence, while the ordinal embedding in SetRank characterizes the rank of a document w.r.t. a ranking algorithm being applied to a document set. Therefore, each word inputted to Transformer has only one position embedding while each document in SetRank may have multiple ordinal embeddings.

4.3 Document Encoding with (Induced) Multi-head Self Attention Block

The key of SetRank is the encoding component which takes the document representations $X^0 = X \in \mathbb{R}^{N \times E}$ as input, and jointly encoding these documents as their internal codes $X^{N_b} \in \mathbb{R}^{N \times E}$. The encoding component is a stack of N_b multi-head self attention blocks (MSAB) (or induced multi-head self attention blocks (IMSAB)) with identical structures. Each MSAB (or IMSAB) receives a set of N vectors (packed as a matrix X^t , $0 \leq t \leq N_b - 1$), processes the vectors with the (induced) multi-head self attention layer

followed by a layer normalization, and then through a row-wise feed-forward network (rFF) layer followed by another layer normalization. Finally, it sends the vectors X^{t+1} as output to the next layer of MSAB (or IMSAB).

Next, we introduce the details of MSAB and IMSAB.

4.3.1 Multi-head self attention block (MSAB). MSAB is based on the attention mechanism in deep neural networks. As have shown in [35], the attention function can be formalized as a scaled dot-product attention with three inputs:

$$\text{Attn}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{E}}\right)V, \quad (3)$$

where $Q \in \mathbb{R}^{N_q \times E}$ denotes the attention query matrix⁰, $K \in \mathbb{R}^{N_k \times E}$ the key matrix, and $V \in \mathbb{R}^{N_v \times E}$ the value matrix. N_q , N_k , and E denote the number of attention query, keys/values vectors, the dimensions of the representation, respectively. The attention mechanism can be explained as: for each attention query vector in Q , it first computes the dot products of the attention query with all keys, aiming to evaluate the similarity between the attention query and each key. Then, it is divided each by \sqrt{E} , and applies a softmax function to obtain the weights on the values. Finally, the new representation of the attention query vector is calculated as weighed sum of values.

To make the attention function more flexible, additional multi-head strategy is usually combined with the attention mechanism. That is, instead of computing a direct attention function, multi-head strategy first projects the inputs Q, K, V into h different spaces, each has the dimensionality of $\hat{E} = E/h$ and projects with the attention function $\text{Attn}(\cdot, \cdot, \cdot)$ defined in Equation (3):

$$\text{Multihead}(Q, K, V) = \text{concat}\left(\left[\text{Attn}(QW_i^Q, KW_i^K, VW_i^V)\right]_{i=1}^h\right) \quad (4)$$

where $W_i^Q, W_i^K, W_i^V \in \mathbb{R}^{E \times \hat{E}}$ reshape the input matrices, and the output $\text{Multihead}(Q, K, V)$ has the same shape with Q .

Based on the Multihead function in Equation (4), the multi-head attention block (MAB) is defined as

$$\begin{aligned} \text{MAB}(Q, K, V) &= \text{LayerNorm}(B + \text{rFF}(B)), \\ \text{where } B &= \text{LayerNorm}(Q + \text{Multihead}(Q, K, V)) \end{aligned} \quad (5)$$

where $\text{rFF}(\cdot)$ is a row-wise feedforward layer, and $\text{LayerNorm}(\cdot)$ is layer normalization [3]. The MAB is an adaptation of the encoder block of the Transformer [35] without positional encoding and dropout. Ideally, K and V can be treated as the index and context of a dictionary, while Q is the information need of the user. For simplicity, we set $K = V$, that means, the index is the context itself.

Finally, the self-attention in MSAB is a special case of MAB when $Q = K = V$, as shown in Figure 2(a). In MSAB, the model represents each document using all the document in the input set:

$$\text{MSAB}(X) = \text{MAB}(X, X, X). \quad (6)$$

4.3.2 Induced Multi-head Self Attention Block (IMSAB). One problem with MSAB is its sensitivity to the size of input set, as will be shown in our experimental studies. The MSAB trained on the data that each query has N_0 labeled documents does not work well on the test time if the queries are associated with $N_1 > N_0$

⁰Note that the “query” in the attention function and IR are with different meanings.

documents. In most cases, N_0 is relatively small because of the expensive and time-consuming relevance label annotation. In real applications, however, it is not nature and realistic to limit the the number of documents that a query can retrieve. Previous studies showed that other multivariate scoring functions also have similar limitations [2].

To address the problem, we employ the induced multi-head self attention [27], or IMSAB, for enabling the model to process the inputted document set with any size. As shown in Figure 2(b), IMSAB can be spilt into two steps. First, it constructs M fake attention query vectors, denoted as $\mathbf{I} \in \mathbb{R}^{M \times E}$, to extract information from original keys/values $\mathbf{X} \in \mathbb{R}^{N \times E}$. Usually, $M < N$ and the M vectors in \mathbf{I} can be viewed as the index of cluster centers for keys/values. The result $\mathbf{H} \in \mathbb{R}^{M \times E}$ is the M cluster centers for keys/values. Then, the original documents $\mathbf{X} \in \mathbb{R}^{N \times E}$ find their contextual representations based on the M cluster centers $\mathbf{H} \in \mathbb{R}^{M \times E}$. The new representations generated by $\text{IMSAB}_M(\mathbf{X})$ have the same shape with the original \mathbf{X} and can be formally written as:

$$\text{IMSAB}_M(\mathbf{X}) = \text{MAB}(\mathbf{X}, \mathbf{H}, \mathbf{H}), \text{ where } \mathbf{H} = \text{MAB}(\mathbf{I}, \mathbf{X}, \mathbf{X}). \quad (7)$$

4.4 Document Ranking

Stacking aforementioned MSAB or IMSAB_M blocks and passing the results to row-wise feedforward neural networks, we achieve two versions of SetRank scoring functions, respectively named as $\text{SetRank}_{\text{MSAB}}$ and $\text{SetRank}_{\text{IMSAB}}$:

$$\begin{aligned} \mathbf{X}_{\text{MSAB}}^{N_b} &= \underbrace{\text{MSAB}(\text{MSAB} \dots (\text{MSAB}(\mathbf{X}^0)))}_{N_b}, \\ \text{SetRank}_{\text{MSAB}}(D) &= \text{rFF}_1(\mathbf{X}_{\text{MSAB}}^{N_b}), \end{aligned} \quad (8)$$

where rFF_1 is a row-wise feed-forward network that projects each document representation into one real value as the corresponding ranking score. Similarly, $\text{SetRank}_{\text{IMSAB}}$ can be obtained by replacing MSAB to IMSAB. The final document rankings, thus, can be achieved by sorting the documents according to the scores:

$$\hat{\pi}_{\text{MSAB}} = \text{sort} \circ \text{SetRank}_{\text{MSAB}}(D). \quad (9)$$

4.5 Model Training

Following the practices in [1], we adopt an attention rank loss function for model training. The basic idea is to measure the distance between an attention distribution generated by the ranking scores and that of generated by the relevance judgments.

Given a labeled query $\psi_q = \{D = \{d_i\}, y = \{y_i\} | 1 \leq i \leq N\}$, where y_i is the relevance label and denotes the information gain of document d_i for query q . The optimal attention allocation strategy in terms of relevance labels for document $d_i (i = 1, \dots, N)$ is:

$$a_i^y = \frac{\tau(y_i)}{\sum_{d_k \in D} \tau(y_k)}, \quad \tau(x) = \begin{cases} \exp(x) & x > 0 \\ 0 & \text{otherwise} \end{cases}$$

Similarly, given the predicted ranking scores $\{s_1, \dots, s_N\} = \text{SetRank}(D)$, the attention w.r.t. the predicted scores for document $d_i (i = 1, \dots, N)$ is:

$$a_i^s = \exp(s_i) / \sum_{d_k \in D} \exp(s_k).$$

Therefore, the list-wise cross entropy loss function can be constructed on the basis of these two attention distributions:

$$\mathcal{L} = \sum_{d_i \in D} a_i^y \log a_i^s + (1 - a_i^y) \log(1 - a_i^s).$$

Note that SetRank is insensitive to the input set size, especially for SetRank based on IMSAB. In the document representation step, ordinal embeddings may be involved to utilize the initial rankings. However, it stop the model from working on the input set size larger than the largest set size in the training data, because the larger ordinal embeddings are unavailable. To address the issue, we proposed a relative ordinal embedding sampling strategy during the model training. First, a max set size N_{\max} (could be a very large number) is chosen and the length of ordinal embeddings is also set to N_{\max} . Then, for each training query, we sample the start number s , constructing a sequence of index numbers $[s, s+1, \dots, s+N-1]$ are used to replace the original document ranking indexes $[1, 2, \dots, N]$. With this trick, all of positions less than N_{\max} can be trained and have their ordinal embeddings. Finally, during the test phase, the ordinal embedding index larger than N and less than N_{\max} is available, so that we can process ranking list smaller than N_{\max} .

4.6 Theoretical Analysis

The existing learning-to-rank algorithms either use uni-variate ranking function that assigns relevance scores to the retrieved documents independently, or use multivariate ranking functions defined on document sequence in which the output is sensitive to the input order of the documents. In contrast, SetRank tries to learn a permutation-invariant ranking model defined on document sets of any size. We show that the RankSet scoring functions defined in Equation (8) are permutation equivariant, proved in Section A.

PROPOSITION 1. *The multivariate scoring functions $\text{SetRank}_{\text{MSAB}}$ and $\text{SetRank}_{\text{IMSAB}}$ are permutation equivariant.*

PROPERTY 1. *The multi-head attention block is permutation equivariant.*

PROPERTY 2. *The induced multi-head self attention block is permutation equivariant.*

The property can be proved in a similar way. Also, Zaheer et al. [41] showed that the composition of permutation equivariant functions is also permutation equivariant. Thus, it is easy to show that $\text{SetRank}_{\text{MSAB}}$ and $\text{SetRank}_{\text{IMSAB}}$ are permutation equivariant.

Since the final outputted document ranking is achieved by sorting with the scores, it is obvious that the two variations of SetRank model shown in Equation (9) are permutation-invariant.

5 EXPERIMENTS

In this section, we describe our experimental settings and results on three publicly available large scale benchmarks.

5.1 Experiment Settings

5.1.1 Datasets. The experiments were conducted on three public available learning-to-rank datasets: Istella LETOR dataset (Istella)¹, Microsoft LETOR 30K (MSLR30K)², and Yahoo! LETOR challenge

¹<http://blog.istella.it/istella-learning-to-rank-dataset/>

²<http://research.microsoft.com/en-us/projects/mslr/>

Table 1: Dataset statistics.

Property	Istella	MSLR 30K	Yahoo!
#features	220	136	700
#queries in training	20,317	18,919	19,944
#queries in validation	2,902	6,306	2,994
Total #docs in train	7,325,625	2,270,296	473,134
#queries in test	9,799	6,306	6,983
Total #docs in test	3,129,004	753,611	165,660
Avg. #docs per query in test	319.31	119.51	23.72

set1 (Yahoo!)³. All these three datasets contain queries and documents sampled from real search engines. Each query-document pair was labeled by human annotators with 5-level relevance judgments ranging from 0 (irrelevant) to 4 (perfectly relevant). As for features, MSLR30K represents each query-document pair with ranking features extracted by human exports; Yahoo! and Istella represent each query-document pair with the actually feature vectors used in the online systems. In the experiments, we conducted 5-fold cross validation on MSLR30K, and used the predefined train/valid/test split in Yahoo! and Istella. More characteristics of the three datasets are listed in Table 1.

5.1.2 Baselines and Evaluation Measures. Learning-to-rank models based on the uni-variate scoring functions were used as the baselines in the experiments. **RANKSVM** [23]: A classic pairwise learning-to-rank model built on the basis of SVM. **RANKBOOST** [15]: Another classic pairwise learning-to-rank model based on boosting; **MART** [16]: A pointwise learning-to-rank algorithm based on gradient boosting regression trees; **LAMBDA MART** [8]: An extension of MART that directly optimize NDCG in training. LambdaMART has been widely used in commercial search engines. The state-of-the-art deep learning-to-rank models based on multivariate scoring functions were also adopted as the baselines. **DLCM** [1]: A deep learning-to-rank model that use local context information to re-rank a list of documents. In DLCM, RNN is used to encode the initial document list. **GSF** [2]: A groupwise scoring model that ranks documents based on group comparisons.

For RankSVM, we used the implementations in SVM^{rank} ⁴; For RankBoost, Mart, and LambdaMART, we used the implementations in *RankLib*⁵. For DLCM and GSF, we used the implementations and settings reported in Github repositories: *DLCM*⁶ and *TF-Ranking*⁷.

All models were tuned and selected based on their performances on the validation set in terms of NDCG@10. Following the practices in [1], we set the RankSVM parameter $C = 200$; the training rounds of RankBoost as 1000. For tree-based models, the number of trees is set to 1000 and the leaf number 20.

Normalized Discounted Cumulative Gain (NDCG) [21] is used to evaluate the performances of the ranking models. To show the model performances at different positions, we reported the NDCG values at the ranks of 1, 3, 5, and 10.

³<http://learningtorankchallenge.yahoo.com>

⁴https://www.cs.cornell.edu/people/tj/svm_light/svm_rank.html

⁵<https://sourceforge.net/p/lemur/wiki/RankLib/>

⁶<https://github.com/QingyaoAi/Deep-Listwise-Context-Model-for-Ranking-Refinement>

⁷https://github.com/tensorflow/ranking/tree/master/tensorflow_ranking

5.1.3 SetRank Settings. We tested four variations of the SetRank, **SetRank_{MSAB}**/**SetRank_{IMSAB}**: using stacked MSAB/IMSAB for documents encoding, without ordinal embeddings; **SetRank_{MSAB}^{init}**/**SetRank_{IMSAB}^{init}**: using stacked MSAB/IMSAB for documents encoding, and ordinal embeddings based on the initial document ranking generated by LambdaMART; In all of the experiments, the SetRank models were configured to have 6 stacked blocks, each has 256 hidden units and 8 heads. In order to fit the hidden dimension of the MSAB or IMSAB, we further appended a rFF_{256} on $\phi(q, d_i)$ to reshape the dimension of d_i to 256. For SetRank_{IMSAB} and SetRank_{IMSAB}^{init}, the induced dimension M is set to 20. The gradient descent method Adam [25] with learning rate 0.001 is used to optimize the objective. Following the practices in [1], in all of the experiments we retrieved top 40 documents per query using LambdaMART. The ranking models were trained on these retrieved documents and then used to re-rank the documents. The initial document rankings using in DLCM, SetRank_{MSAB}^{init} and SetRank_{IMSAB}^{init} are generated by LambdaMART. For fair comparison, we also evaluated DLCM without the initial rankings, denotes as $DLCM^{w/o init}$.

The source code and experiments of SetRank can be found at <https://github.com/pl8787/SetRank>.

5.2 Experimental Results

Table 2 lists the experimental results on the three datasets: Istella dataset, MSLR30K dataset, and Yahoo! dataset, respectively. For convenient comparisons, we categorized the baselines and the four versions of SetRank into three groups: models based on uni-variate scoring functions (i.e., RankSVM, RankBoost, Mart, and LambdaMart), models based on multivariate scoring functions and without initial rankings (i.e., $DLCM^{w/o init}$, GSF, SetRank_{MSAB}, and SetRank_{IMSAB}), and models based on multivariate scoring functions and with initial rankings by LambdaMART (i.e., DLCM, SetRank_{MSAB}^{init}, and SetRank_{IMSAB}^{init}).

Comparing the performances of the first group and the second (and the third) groups, we can see that the ranking models based on multivariate scoring functions are in general performed better than those based on uni-variate scoring functions. The multivariate models using LambdaMart as the initial rankings performed significantly better than LambdaMart, though LambdaMart is a very strong baseline. The results indicate that the multivariate scoring functions are superior to uni-variate ones because they captured the contextual information and document interactions during ranking.

Also noted that SetRank outperformed the baselines of DLCM and GSF, in the cases of both with and without the initial rankings. In Istella dataset, without the initial rankings, SetRank_{IMSAB} improved $DLCM^{w/o init}$ about 0.05 points in terms of NDCG@1; with the initial rankings, SetRank_{IMSAB}^{init} improved DLCM about 0.02 points in term of NDCG@1. Similar phenomenon can be observed on other datasets with other evaluation measures. As discussed previously, the RNN in DLCM is very sensitive to the orders the document being inputted [1]. SetRank, on the other hand, learns a permutation-invariant ranking model and thus is more robust.

Comparing the performances of SetRank_{MSAB} and SetRank_{IMSAB}, in most cases, SetRank_{IMSAB} achieved better document rankings. This is because the IMSAB blocks can automatically adapt to input set of any size, which makes the ranking model more robust.

Table 2: Performance comparison of different models on Istella, MSLR30K and Yahoo datasets. Significant performance improvements (paired t-test with p-value ≤ 0.05) over LambdaMart and DLCM are denoted as ‘+’ and ‘†’, respectively. Boldface indicates the best performed results.

(a) Ranking accuracies on Istella LETOR dataset				
NDCG				
Model	@1	@3	@5	@10
RankSVM	0.5269	0.4867	0.5041	0.5529
RankBoost	0.4457	0.3977	0.4097	0.4511
Mart	0.6185	0.5633	0.5801	0.6285
LambdaMart	0.6571	0.5982	0.6118	0.6591
Without initial rankings				
DLCM ^{w/o init}	0.6272	0.5717	0.5848	0.6310
GSF	0.6224	0.5796	0.5968	0.6508
SetRank _{MSAB}	0.6702 ⁺ †	0.6150 ⁺ †	0.6282 ⁺ †	0.6766 ⁺ †
SetRank _{ISAB}	0.6733 ⁺ †	0.6136 ⁺ †	0.6278 ⁺ †	0.6737 ⁺ †
With initial rankings generated by LambdaMart				
DLCM	0.6558	0.6030 ⁺	0.6194 ⁺	0.6680 ⁺
SetRank _{MSAB} ^{init}	0.6745 ⁺ †	0.6201 ⁺ †	0.6350⁺†	0.6819 ⁺ †
SetRank _{ISAB} ^{init}	0.6760⁺†	0.6202⁺†	0.6345 ⁺ †	0.6834⁺†
(b) Ranking accuracies on Microsoft LETOR 30K dataset				
NDCG				
Model	@1	@3	@5	@10
RankSVM	0.3010	0.3180	0.3350	0.3650
RankBoost	0.2788	0.2897	0.3043	0.3339
Mart	0.4436	0.4344	0.4414	0.4633
LambdaMart	0.4570	0.4420	0.4450	0.4640
Without initial rankings				
DLCM ^{w/o init}	0.3985	0.3919	0.4001	0.4245
GSF	0.4129	0.4073	0.4151	0.4374
SetRank _{MSAB}	0.4243	0.4116	0.4177	0.4403
SetRank _{ISAB}	0.4290	0.4166	0.4220	0.4428
With initial rankings generated by LambdaMart				
DLCM	0.4630⁺	0.4450 ⁺	0.4500 ⁺	0.4690 ⁺
SetRank _{MSAB} ^{init}	0.4572	0.4452 ⁺	0.4499 ⁺	0.4692 ⁺
SetRank _{ISAB} ^{init}	0.4591 ⁺	0.4469⁺	0.4515⁺	0.4696⁺
(c) Ranking accuracies on Yahoo! LETOR challenge set1 dataset				
NDCG				
Model	@1	@3	@5	@10
RankSVM	0.6370	0.6500	0.6740	0.7260
RankBoost	0.6293	0.6409	0.6661	0.7159
Mart	0.6830	0.6827	0.7034	0.7469
LambdaMart	0.6770	0.6760	0.6960	0.7380
Without initial rankings				
DLCM ^{w/o init}	0.6693	0.6751	0.6958	0.7391
GSF	0.6429	0.6604	0.6838	0.7316
SetRank _{MSAB}	0.6623	0.6698	0.6911	0.7369
SetRank _{ISAB}	0.6711	0.6760	0.6960	0.7398
With initial rankings generated by LambdaMart				
DLCM	0.6760	0.6810 ⁺	0.6990 ⁺	0.7430 ⁺
SetRank _{MSAB} ^{init}	0.6837⁺†	0.6820 ⁺	0.7009 ⁺	0.7443 ⁺
SetRank _{ISAB} ^{init}	0.6822 ⁺ †	0.6835⁺†	0.7029 ⁺	0.7453 ⁺ †




Figure 3: The performance of models w.r.t. the number reversed document pairs per query in the initial rankings.

5.3 Empirical Analyses

We now investigate the reasons that SetRank outperforms the baseline methods, using the results on Istella dataset as examples.

5.3.1 The effects of permutation invariance. We conducted experiments to show that permutation invariance property helps to learn a ranking model robust to the noises in the initial ranking list. Specifically, we randomly reversed some document pairs in the ranking lists generated by LambdaMart for each query. The processed document lists are then feeded as the initial rankings to DLCM and SetRank^{init}_{ISAB} in the test phase, while both of DLCM and SetRank^{init}_{ISAB} model are trained using the clean ranking list.

Figure 3 illustrates the performance curves in term of NDCG@10 w.r.t. the number reversed document pairs. We can see that, with the increasing number of reverse pairs (i.e., with more noise) in the initial rankings, both SetRank^{init}_{ISAB} and DLCM performed worse. However, much more performance declines are observed with DLCM. This is because DLCM encodes the input documents with an RNN, which is sensitive to the permutations on the input set.

We also illustrate the performance curve of SetRank_{ISAB}, which did not use the ordinal embedding, in Figure 3. We can see that the performances of SetRank_{ISAB} did not change w.r.t. the number of reversed document pairs. The results confirmed the correctness of the theoretical analysis in Proposition 1.

5.3.2 The effects of set size adaptation. In this part, we want to test the effects of ISMAB which enables the set size adaptation ability of SetRank. We first trained and tested the DLCM and SetRank models (without the initial rankings) on the data that each query was associated with N documents ($N = N_0 = N_1 = 40, 240$, or 500). From the corresponding results shown in Table 3 (lines in which “# docs associated per query” in training set and in test set are identical), we can see that the three models, DLCM^{w/o init}, SetRank_{MSAB} and SetRank_{ISAB} have similar performances on different set sizes. It indicates that when the document sets have similar sizes in phase of training and test, the performances of the models tend to be unchanged. Nevertheless, when we trained the DLCM and SetRank models on queries associated with 40 documents ($N_0 = 40$) and tested the models on queries associated with 240 and 500 documents ($N_1 = 240$ or 500). From the corresponding results reported in Table 3 (lines in which “# docs associated per query” in training set is 40, and in test set is 240 or 500), we can see that the performance drop (see the Δ values in the brackets) in terms of NDCG@10 for DLCM^{w/o init} is big, indicating that

Table 3: Performances of the models trained and tested on queries that are associated with different number of documents. The values in the parentheses (Δ) are the performance variations compared with the models trained and tested on set size 40.

Model	# docs associated per query		NDCG@1 (Δ)	NDCG@3 (Δ)	NDCG@5 (Δ)	NDCG@10 (Δ)
	training set (N_0)	test set (N_1)				
DLCM ^{w/o init}	40	40	0.6233	0.5684	0.5825	0.6298
		240	0.5943 (-0.0290)	0.5394 (-0.0290)	0.5518 (-0.0307)	0.5964 (-0.0334)
		500	0.5844 (-0.0389)	0.5300 (-0.0384)	0.5428 (-0.0397)	0.5868 (-0.0430)
	240	240	0.6199	0.5708	0.5836	0.6309
SetRank _{MSAB}	40	40	0.6702	0.6150	0.6282	0.6766
		240	0.6578 (-0.0124)	0.6026 (-0.0124)	0.6155 (-0.0127)	0.6602 (-0.0164)
		500	0.6533 (-0.0169)	0.5969 (-0.0181)	0.6102 (-0.0180)	0.6547 (-0.0219)
	240	240	0.6736	0.6141	0.6295	0.6777
SetRank _{IMSAB}	40	40	0.6733	0.6136	0.6278	0.6737
		240	0.6674 (-0.0059)	0.6104 (-0.0032)	0.6244 (-0.0034)	0.6688 (-0.0049)
		500	0.6665 (-0.0068)	0.6082 (-0.0054)	0.6220 (-0.0058)	0.6662 (-0.0075)
	240	240	0.6699	0.6115	0.6264	0.6750
	500	500	0.6696	0.6117	0.6247	0.6727

Table 4: Performances of SetRank variations that involve zero, one, or multiple initial rankings.

Model	NDCG			
	@1	@3	@5	@10
SetRank _{IMSAB}	0.6733	0.6136	0.6278	0.6737
SetRank _{IMSAB} ^{init}	0.6760	0.6202	0.6345	0.6834
SetRank _{IMSAB} ^{multi-init}	0.6744	0.6211	0.6365	0.6860

the sequential model DLCM^{w/o init} is sensitive to the test set size. For example, when the models were trained on queries associated with 40 documents while tested on queries with 240 documents, the performance drop for DLCM^{w/o init} is 0.0334, which is double compared to SetRank_{MSAB} and six times compared to SetRank_{IMSAB}.

Additionally, note that the performance drop of SetRank_{MSAB} is much severe than that of for SetRank_{IMSAB}, indicating that projecting variance set size to a fixed number of cluster centers as IMSAB does, reduces the sensitivity of the test set size. For example, when the models were trained on queries associated with 40 documents while tested on queries with 240 documents, the performance drop for SetRank_{IMSAB} in terms of NDCG@10 is 0.0049, while the number is 0.0164 for SetRank_{MSAB}.

5.3.3 The effects of initial rankings. One advantage of SetRank is that it can involve zero, one, or multiple initial document rankings as inputs, through adding one or more ordinal embeddings. We conducted experiments to test the effects of the initial rankings in SetRank. Specifically, we tested the performances of SetRank_{IMSAB} with no initial ranking, SetRank_{IMSAB}^{init} with one initial ranking based on LambdaMart, and SetRank_{IMSAB}^{multi-init} with four initial rankings based on LambdaMart, Mart, RankSVM, and RankBoost. Table 4 reports the performances of these three SetRank variations in terms of NDCG at the positions of 1, 3, 5, and 10. From the results, we found that (1) SetRank_{IMSAB}^{init} with one initial ranking performed better than SetRank_{IMSAB} which has no initial ranking, indicating that initial rankings do help to improve the performances;

(2) SetRank_{IMSAB}^{multi-init} with 4 initial rankings performed a little bit better than SetRank_{IMSAB}^{init} which used only one initial ranking, indicating that multiple rankings provide limited helps in improving the ranking accuracy. One possible reason is that LambdaMART has provided significantly better document rankings comparing to other ranking models, which makes the information obtained from other ranking models less useful in our experiments.

6 CONCLUSION

In this paper, we proposed a novel learning-to-rank model for information retrieval, referred to as SetRank. In contrast to existing models, the scoring function in SetRank is designed as a multivariate mapping from a document set to a permutation, and satisfy two critical requirements: cross-document interactions and permutation invariance. Self-attention mechanism in Transformer is employed to implement the scoring function. SetRank offers several advantages: efficiently capturing local context information, naturally involving (multiple) initial rankings, robust to input noise, and high accuracy in ranking. Experimental results on three large scale datasets show that SetRank outperformed the traditional learning-to-rank models and state-of-the-art deep ranking models. Analyses showed that the two requirements did help to improve the performance and robustness of SetRank model.

ACKNOWLEDGMENTS

This work was supported by Beijing Academy of Artificial Intelligence (BAAI) under Grants BAAI2020ZJ0303 and BAAI2019ZD0305, the National Natural Science Foundation of China (NSFC) under Grants No. 61773362, 61872338, 61832017, and 61906180, the Youth Innovation Promotion Association CAS under Grants No. 2016102, the Tencent AI Lab Rhino-Bird Focused Research Program (No. JR202033), the National Key R&D Program of China under Grants No. 2016QY02D0405, the Beijing Outstanding Young Scientist Program No. BJJWZYJH012019100020098, Fundamental Research Funds for the Central Universities, and Research Funds of Renmin University of China (2018030246).

REFERENCES

- [1] Qingyao Ai, Keping Bi, Jiafeng Guo, and W Bruce Croft. 2018. Learning a deep listwise context model for ranking refinement. In *Proceedings of the 41st ACM SIGIR*. ACM, 135–144.
- [2] Qingyao Ai, Xuanhui Wang, Nadav Golbandi, Michael Bendersky, and Marc Najork. 2019. Learning groupwise scoring functions using deep neural networks. In *Proceedings of the ACM ICTIR*.
- [3] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. 2016. Layer normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [4] Irwan Bello, Sayali Kulkarni, Sagar Jain, Craig Boutilier, Ed Chi, Elad Eban, Xiyang Luo, Alan Mackey, and Ofer Meshi. 2018. Seq2slate: Re-ranking and slate optimization with rnns. *arXiv preprint arXiv:1810.02019* (2018).
- [5] Sebastian Bruch, Shuguang Han, Michael Bendersky, and Marc Najork. 2020. A Stochastic Treatment of Learning to Rank Scoring Functions. In *Proceedings of the 13th WSDM*. ACM, 61–69.
- [6] Sebastian Bruch, Masrour Zoghi, Michael Bendersky, and Marc Najork. 2019. Revisiting Approximate Metric Optimization in the Age of Deep Neural Networks. In *Proceedings of the 42nd International ACM SIGIR*. ACM, 1241–1244.
- [7] Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. 2005. Learning to rank using gradient descent. In *Proceedings of the 22nd ICML*. ACM, 89–96.
- [8] Christopher JC Burges. 2010. From ranknet to lambdarank to lambdamart: An overview. *Learning* 11 (2010), 23–581.
- [9] Ethem F Can, W Bruce Croft, and R Manmatha. 2014. Incorporating query-specific feedback into learning-to-rank models. In *Proceedings of the 37th ACM SIGIR*. ACM, 1035–1038.
- [10] Zhe Cao, Tao Qin, Tie-Yan Liu, Ming-Feng Tsai, and Hang Li. 2007. Learning to rank: from pairwise approach to listwise approach. In *Proceedings of the 24th ICML*. ACM, 129–136.
- [11] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Āukasz Kaiser. 2018. Universal Transformers. *arXiv:cs.CL/1807.03819*
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics*. 4171–4186.
- [13] Yajuan Duan, Long Jiang, Tao Qin, Ming Zhou, and Heung-Yeung Shum. 2010. An empirical study on learning to rank of tweets. In *Proceedings of the 23rd International Conference on Computational Linguistics*. ACL, 295–303.
- [14] Harrison Edwards and Amos J. Storkey. 2017. Towards a Neural Statistician. In *5th International Conference on Learning Representations, ICLR 2017*.
- [15] Yoav Freund, Raj Iyer, Robert E Schapire, and Yoram Singer. 2003. An efficient boosting algorithm for combining preferences. *Journal of machine learning research* 4, Nov (2003), 933–969.
- [16] Jerome H Friedman. 2001. Greedy function approximation: a gradient boosting machine. *Annals of statistics* (2001), 1189–1232.
- [17] Norbert Fuhr. 2008. A probability ranking principle for interactive information retrieval. *Information Retrieval* 11, 3 (2008), 251–265.
- [18] Jiafeng Guo, Yixing Fan, Qingyao Ai, and W Bruce Croft. 2016. A deep relevance matching model for ad-hoc retrieval. In *Proceedings of the 25th ACM CIKM*. ACM, 55–64.
- [19] Baotian Hu, Zhengdong Lu, Hang Li, and Qingcai Chen. 2014. Convolutional neural network architectures for matching natural language sentences. In *NIPS*. 2042–2050.
- [20] Po-Sen Huang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Acero, and Larry Heck. 2013. Learning deep structured semantic models for web search using clickthrough data. In *Proceedings of the 22nd ACM CIKM*. ACM, 2333–2338.
- [21] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [22] Ray Jiang, Sven Gowal, Timothy A Mann, and Danilo J Rezende. 2018. Beyond greedy ranking: Slate optimization via List-CVAE. *arXiv preprint arXiv:1803.01682* (2018).
- [23] Thorsten Joachims. 2006. Training linear SVMs in linear time. In *Proceedings of the 12th ACM SIGKDD*. ACM, 217–226.
- [24] Thorsten Joachims, Laura A Granka, Bing Pan, Helene Hembrooke, and Geri Gay. 2005. Accurately interpreting clickthrough data as implicit feedback. In *SIGIR*, Vol. 5. 154–161.
- [25] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [26] Victor Lavrenko and W Bruce Croft. 2017. Relevance-based language models. In *ACM SIGIR Forum*, Vol. 51. ACM, 260–267.
- [27] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. 2019. Set Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In *Proceedings of the 36th ICML*. Long Beach, California, USA, 3744–3753.
- [28] Tie-Yan Liu. 2009. Learning to rank for information retrieval. *Foundations and Trends in Information Retrieval* 3, 3 (2009), 225–331.
- [29] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Shengxian Wan, and Xueqi Cheng. 2016. Text matching as image recognition. In *Thirtieth AAAI Conference on Artificial Intelligence*.
- [30] Liang Pang, Yanyan Lan, Jiafeng Guo, Jun Xu, Jingfang Xu, and Xueqi Cheng. 2017. DeepRank: A new deep architecture for relevance ranking in information retrieval. In *Proceedings of the 2017 ACM CIKM*. ACM, 257–266.
- [31] Rama Kumar Pasumarthi, Xuanhui Wang, Michael Bendersky, and Marc Najork. 2019. Self-Attentive Document Interaction Networks for Permutation Equivariant Ranking. *arXiv:cs.IR/1910.09676*
- [32] Changhua Pei, Yi Zhang, Yongfeng Zhang, Fei Sun, Xiao Lin, Hanxiao Sun, Jian Wu, Peng Jiang, Junfeng Ge, Wenwu Ou, and Dan Pei. 2019. Personalized Re-ranking for Recommendation. In *Proceedings of the 13th ACM Conference on Recommender Systems*. 3–11.
- [33] Stephen E Robertson. 1977. The probability ranking principle in IR. *Journal of documentation* 33, 4 (1977), 294–304.
- [34] Falk Scholer, Andrew Turpin, and Mark Sanderson. 2011. Quantifying test collection quality based on the consistency of relevance judgements. In *Proceedings of the 34th ACM SIGIR*. ACM, 1063–1072.
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*. 5998–6008.
- [36] Shengxian Wan, Yanyan Lan, Jun Xu, Jiafeng Guo, Liang Pang, and Xueqi Cheng. 2016. Match-SRNN: modeling the recursive matching structure with spatial RNN. In *Proceedings of the 25th IJCAI*. 2922–2928.
- [37] Fen Xia, Tie-Yan Liu, Jue Wang, Wensheng Zhang, and Hang Li. 2008. Listwise approach to learning to rank: theory and algorithm. In *Proceedings of the 25th ICML*. ACM, 1192–1199.
- [38] Liu Yang, Qingyao Ai, Damiano Spina, Ruey-Cheng Chen, Liang Pang, W Bruce Croft, Jiafeng Guo, and Falk Scholer. 2016. Beyond Factoid QA: Effective Methods for Non-factoid Answer Sentence Retrieval. In *Proceedings of the ECIR*. Springer, 115–128.
- [39] Ziyang Yang. 2017. Relevance Judgments: Preferences, Scores and Ties. In *Proceedings of the 40th ACM SIGIR*. ACM, 1373–1373.
- [40] Emine Yilmaz, Manisha Verma, Nick Craswell, Filip Radlinski, and Peter Bailey. 2014. Relevance and effort: An analysis of document utility. In *Proceedings of the 23rd ACM CIKM*. ACM, 91–100.
- [41] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Ruslan R Salakhutdinov, and Alexander J Smola. 2017. Deep sets. In *NeurIPS*. 3391–3401.
- [42] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. 2019. Self-Attention Generative Adversarial Networks. In *ICML*. 7354–7363.
- [43] Yongfeng Zhang, Xu Chen, Qingyao Ai, Liu Yang, and W Bruce Croft. 2018. Towards conversational search and recommendation: System ask, user respond. In *Proceedings of the 27th ACM CIKM*. ACM, 177–186.

A PROOF TO PROPOSITION 1

Note that the multi-head parameters $\mathbf{W}_i^Q, \mathbf{W}_i^K, \mathbf{W}_i^V$, the row-wise feedforward function $\text{rFF}(\cdot)$, and layer normalization $\text{LayerNorm}(\cdot)$ are all element wise functions, which do not affect the permutation-equivariant property. Therefore, we only need to consider the permutation-equivariant of the self-attention function:

$$\text{SelfAttn}(\mathbf{X}) = \text{Attn}(\mathbf{X}, \mathbf{X}, \mathbf{X}) = \text{softmax} \left(\frac{\mathbf{X}\mathbf{X}^T}{\sqrt{E}} \right) \mathbf{X}.$$

PROOF. Let $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N\}$. For simplicity and with no loss of generality we discard the scalar factor \sqrt{E} , we have

$$\begin{aligned} \text{SelfAttn}(\mathbf{X}) &= \text{softmax} \left(\mathbf{X}\mathbf{X}^T \right) \mathbf{X} = \text{softmax} \left(\left[\mathbf{x}_i \mathbf{x}_j^T \right]_{ij} \right) \mathbf{X} \\ &= \left[\frac{e^{\mathbf{x}_i \mathbf{x}_j}}{\sum_k e^{\mathbf{x}_i \mathbf{x}_k}} \right]_{ij} \mathbf{X} = \left[\frac{\sum_j e^{\mathbf{x}_i \mathbf{x}_j} \cdot \mathbf{x}_j}{\sum_k e^{\mathbf{x}_i \mathbf{x}_k}} \right]_i, \end{aligned}$$

For any permutation $\pi \in \Pi_N$, since sum is permutation-invariant,

$$\begin{aligned} \text{SelfAttn}(\pi \mathbf{X}) &= \text{SelfAttn} \left(\left[\mathbf{x}_{\pi(i)} \right]_i \right) = \left[\frac{\sum_j e^{\mathbf{x}_{\pi(i)} \mathbf{x}_{\pi(j)}} \cdot \mathbf{x}_{\pi(j)}}{\sum_k e^{\mathbf{x}_{\pi(i)} \mathbf{x}_{\pi(k)}}} \right]_{\pi(i)} \\ &= \left[\frac{\sum_j e^{\mathbf{x}_i \mathbf{x}_j} \cdot \mathbf{x}_j}{\sum_k e^{\mathbf{x}_i \mathbf{x}_k}} \right]_{\pi(i)} = \pi \left(\text{SelfAttn}(\mathbf{X}) \right). \end{aligned}$$

□