

# COLTR: Semi-supervised Learning to Rank with Co-training and Over-parameterization for Web Search

Yuchen Li, Haoyi Xiong, *Senior Member, IEEE*, Qingzhong Wang, Linghe Kong, *Senior Member, IEEE*, Hao Liu, Haifang Li, Jiang Bian, Shuaiqiang Wang, Guihai Chen, *Fellow, IEEE*, Dejing Dou, *Senior Member, IEEE*, Dawei Yin

**Abstract**—While *learning to rank* (LTR) has been widely used in web search to prioritize most relevant webpages among the retrieved contents subject to the input queries, the traditional LTR models fail to deliver decent performance due to two main reasons: 1) the lack of well-annotated query-webpage pairs with ranking scores to cover search queries of various popularity, and 2) ill-trained models based on a limited number of training samples with poor generalization performance. To improve the performance of LTR models, tremendous efforts have been done from above two aspects, such as enlarging training sets with pseudo-labels of ranking scores by self-training, or refining the features used for LTR through feature extraction and dimension reduction. Though LTR performance has been marginally increased, we still believe these methods could be further improved in the newly-fashioned “interpolating regime”. Specifically, instead of lowering the number of features used for LTR models, our work proposes to transform original data with random Fourier feature, so as to over-parameterize the downstream LTR models (e.g., GBRank or LightGBM) with features in ultra-high dimensionality and achieve superb generalization performance. Furthermore, rather than self-training with pseudo-labels produced by the same LTR model in a “self-tuned” fashion, the proposed method incorporates the diversity of prediction results between the listwise and pointwise LTR models while co-training both models with a cyclic labeling-prediction pipeline in a “ping-pong” manner. We deploy the proposed *Co-trained and Over-parameterized LTR* system **COLTR** at Baidu search and evaluate **COLTR** with a large number of baseline methods. The results show that **COLTR** could achieve  $\Delta NDCG_4=3.64\%\sim 4.92\%$ , compared to baselines, under various ratios of labeled samples. We also conduct a 7-day A/B Test using the realistic web traffics of Baidu Search, where we can still observe significant performance improvement around  $\Delta NDCG_4=0.17\%\sim 0.92\%$  in real-world applications. **COLTR** performs consistently both in online and offline experiments.

**Index Terms**—Learning to Rank, Semi-supervised Learning, Over-parameterization

## 1 INTRODUCTION

INCORPORATING with billions of population and trillions of web content, search service like Google, Baidu, and Bing satisfy the daily searching needs of users. The rapid growth of web information and internet users surge the needs of web search engine. Nowadays, orchestrating with trillions of webpages archived and indexed for search, a large-scale industrial search engine serves hundreds of millions of daily active users and handles billions of queries

per day. In addition to users, webpages, and computing resources, to improve the web service, a great number of most advanced technologies have been invented, ranging from pre-trained language models for content/query understanding [1], [2], [3] and ranking [4], domain-specific recommender systems for personalized recommendation [5], online query-Ads matching for sponsored search [6], [7] and advanced infrastructures with software/hardware co-design [8], [9] for handling web-scale traffic of online search.

With a query (e.g., a string of texts) input by the user, a search engine needs first to extract keywords/phrases from the query and recognize the user’s intention. Given the extracted keywords or phrases, the search engine evaluates the relevance between the query and webpages, and after that, retrieves a list of webpages that are relevant from a database of trillions of webpages. Further, the search engine ranks the retrieved webpages based on their contents and click-through rates. In the response to the query, the search engine tops the most relevant webpages. To optimize the user experience of search, ranking the retrieved contents is a key step, where *Learning to Rank* (LTR) plays a critical role.

To achieve high accuracy for LTR, there needs to collect ultra-large training datasets and train LTR models for ranking. Specifically, given a large volume of queries with varying popularity (from highly frequent to infrequent queries), the search engine needs to first label the rank of every relevant webpage for every query and then train the LTR models through supervised learning. However, it

*This work was supported in part by National Key R&D Program of China (No. 2021ZD0110303), NSFC grant 62141220, 61972253, U1908212, 62172276, 61972254, the Program for Professor of Special Appointment (Eastern Scholar) at Shanghai Institutions of Higher Learning, Shanghai Science and Technology Development Funds 23YF1420500, Open Research Projects of Zhejiang Lab No. 2022NLOAB01. (Corresponding author: Linghe Kong and Haoyi Xiong.)*

- Yuchen Li, Linghe Kong, and Guihai Chen are with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai 200240, China. E-mail: yuchenli@sjtu.edu.cn; linghe.kong@sjtu.edu.cn; gchen@cs.sjtu.edu.cn.
- Haoyi Xiong, Qingzhong Wang, Haifang Li, Jiang Bian, Shuaiqiang Wang, Dejing Dou, and Dawei Yin are with Baidu, Inc., Beijing 100085, China. E-mail: haoyi.xiong.fr@ieee.org; wangqingzhong@baidu.com; li-haifang@baidu.com; jiangbian03@gmail.com; shqiang.wang@gmail.com; yindawei@acm.org.
- Hao Liu is with Thrust of Artificial Intelligence, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou, Guangdong Province 510000, China, and also with the Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Hong Kong SAR 999077, China. E-mail: liuh@ust.hk.
- The work was done when Yuchen Li did his internship in Baidu, Inc., Beijing, China.

is extremely expensive and time-consuming to label the ranks of relevant webpages for every query [10]. To address this issue, it frequently needs to incorporate both labeled and unlabeled query-webpage pairs to train LTR models in a semi-supervised learning setting. However, semi-supervised LTR at web-scale is not easy, as it might be necessary to leverage trillions of webpages under billions of queries without ranking scores to improve LTR models training based on an extremely small number of labeled samples.

A simple way to enable semi-supervised learning is *self-training* [11], [12], where the machine learning model was (1) first trained with the labeled samples, then (2) predicted over unlabeled data to obtain pseudo-labels. Later, (3) the self-trainer paired unlabeled samples with pseudo-labels and re-trained the model incorporating with both labeled and pseudo-labeled samples. Above steps (2) and (3) could recycle in a self-tuned manner until achieving the best validation performance. Compared to other semi-supervised learning algorithms, such as graph-based or metric learning based approach [13], [12], [14], [15], the *self-training* mechanism is a scalable yet effective method blessed by its low-complexity nature. For example, it has been used to boost the performance of ImageNet classification tasks through incorporating large-scale unlabeled images [16]. The performance of self-training however would be bottlenecked for LTR tasks by issues as follows:

- *Low-Cost Representation for LTR*. For web search in practices, LTR tasks are usually based on statistical learners [17], [18], [19], [20] with a small number of “*hand-brewed*” features (e.g., from several dozens to hundreds of features) [21], [22], [23]. Even when deep models, such as pre-trained language models, are incorporated for feature extraction from raw queries/webpages, LTR models are usually not trained with the deep feature extractors in an end-to-end manner for industry practices. Actually, a post-hoc approach connecting deep representations with RankSVM [24]/GBRank [25] is more preferred [26]. It is because language features would rarely change in a short period but the internet interests (news, celebrities, etc.) shift in an extremely fast manner. Thus, web-scale search engines request to re-train LTR models using the most recent collection of queries/webpages and update the online LTR models frequently, with a “*just-in-time*” solution (but not to re-train language models). In this way, there needs to model *Learning to Rank* (LTR) to a statistical learning task with low-cost representations.
- *Diversifying Noisy Supervision Signals*. Yet another problem of self-training is over-fitting to inaccurate pseudo-labels (e.g., noisy supervision signals), as the LTR model learns from a set of pseudo-labels derived from the (inaccurate) prediction results of an LTR model trained in the previous round. One way to solve the problem is to co-train the LTR model with multiple classifiers [27], [28], [29], so as to incorporate the diversity of prediction outputs from multiple classifiers in an ensemble learning fashion [30], [31]. It has been found that strong learners could be trained with limited labeled samples by making weak learners (producing inaccurate but

diverse prediction results) teach each other [32]. In this way, there needs a way to co-train multiple LTR models while making their prediction results diverse, using the same set of labeled/unlabeled samples.

To scale-up semi-supervised learning for LTR at web-scale, we propose **COLTR** — *Co-trained and Over-parameterized LTR models*, where we solve aforementioned two technical issues with *random Fourier feature (RFF) based over-parameterization* and *multi-loss co-training* strategies respectively. Specifically, inspired by the recently observed phenomenon “*double descent*” of generalization performance [33] with increasing complexity of models, **COLTR** adopts feature-wise “*double descent*” and leverages random Fourier features (RFF) to extend the dimensions of features for LTR data (e.g., query-webpage pairs). With RFF transformed samples, **COLTR** could over-parameterize LTR models so as to enable the representation learning in the so-called interpolating regime [34] with superb performance improvement. Furthermore, **COLTR** co-trains dual LTR models with *listwise* and *pointwise* losses respectively, in a loop of multiple rounds. Specifically, **COLTR** first trains an LTR model based on the *listwise* loss using both labeled/unlabeled query-webpage pairs in a self-training manner, and then generates pseudo-labels for unlabeled samples to train another LTR model based on the *pointwise* loss. Later, **COLTR** makes these two LTR models teach each other in multiple rounds, with pseudo-labels updated. In summary, this work makes contributions as follows:

- We study the problem of semi-supervised *learning to rank* in the context of web-scale search, where we particularly focus on the technical challenges of constructing *low-cost statistical representations for LTR* and *diversifying noisy supervision signals*. To the best of our knowledge, this work is the first to study semi-supervised training for LTR models with labeled/unlabeled query-webpage pairs by addressing the mathematical phenomenon of interpolating [33], [34] in LTR tasks and the diversity of LTR models trained with various losses (e.g., pointwise, pairwise, and listwise) functions.
- We design and implement **COLTR**, incorporating both labeled or unlabeled query-webpage pairs for training LTR models in a semi-supervised manner. Given a query and indexed webpages, **COLTR** relies on a pre-trained language models based retrieval method to pickup the candidates of webpages for ranking, constructs their LTR features using the outputs of language models, and predicts the order of webpages among the candidates using LTR models. To train the LTR models, **COLTR** consists of three steps: (1) *RFF-based over-parameterization* that pushes the limits of representation learning to interpolating regime [34], (2) *Listwise-based Self-training* that initializes pseudo-labels of unlabeled samples using the predictions of an LTR model trained by the listwise loss, and (3) *Multi-Loss Co-training for LTR* that makes pointwise and listwise models learn from each other for multiple rounds with pseudo-labels updated by predictions. Note that **COLTR** restores the RFF-transformed representation of LTR features for either queries or webpages as the immediate results of ranking and accelerates the inference procedure for

online ranking accordingly.

- We deploy COLTR at Baidu Search<sup>1</sup> and evaluate the proposed algorithm using both offline experiments and online A/B Test in comparison with baseline algorithms. The experiment results show that, compared to the state of the art in webpage ranking, COLTR could achieve  $\Delta NDCG_4=3.64\%\sim 4.92\%$  in offline experiments and  $\Delta NDCG_4=0.17\%\sim 0.92\%$  in online A/B tests under fair comparisons. Ablation studies further confirm the effectiveness of *RFF-based model over-parameterization* and *multi-loss co-training* for LTR.

Note that we focus on low-complexity strategies for semi-supervised LTR that can scale-up on real-world traffics, thus advanced methods with higher complexity are not in the scope of our study. Furthermore, please be advised that COLTR is a semi-supervised LTR component in Baidu Search. Experiment results reported in this study are based on A/B tests with the *status quo* of Baidu Search, which has already secured excellent LTR performance.

## 2 RELATED WORK

In this section, we review and discuss related works from the following three aspects: (1) *Learning to Rank*, (2) *Over-parameterization Method* and (3) *Semi-supervised Learning*.

### 2.1 Learning to Rank

To improve user experience in terms of searching, ranking the retrieved contents is a key step, where the LTR model plays a critical role. According to the loss function, we could categorize the LTR models into three families: pointwise [35], [36], pairwise [17], [24] and listwise [37], [38]. The pointwise model (e.g., McRank [36]) formulates the ranking problem into regression tasks to fit the labels of query-webpage pairs. The pairwise model (e.g., RankNet[39]) converts two documents into a document pair and recasts the LTR tasks as binary classification problems. It pays more attention to finding the best one in each document pair. The listwise model (e.g., Sofrank [37]) treats the whole document list as a sample and directly optimizes the evaluation metrics, such as the utilized metric in this work, i.e., NDCG [40], [41], [42]. In general, listwise models can gain the best performance among the three LTR methods. However, pairwise and listwise models are generally deployed in real-world applications for being easy to apply and having less computational complexity. In recent years, deep models have been applied to LTR tasks through end-to-end minimization of various ranking loss functions (e.g., pointwise, listwise, pairwise, and their variants/surrogates) [41], [43], [44], [42]. Moreover, the need for extensive data annotation has prompted both academia and industry to focus on recommendation or LTR tasks using user feedback with deep learning technical [45], [46], [47]. Unbiased LTR is proposed to mitigate biases in feedback and achieves significant performance in many scenarios [48], [49], [50]. *In our work, COLTR considers the divergence between the prediction results of listwise and pointwise models and incorporates such divergence to improve co-training.*

1. <https://www.baidu.com/>

### 2.2 Over-parameterization Method

Recently, the methods of balancing under-parameterization and over-parameterization have attracted growing research interests [33]. [51] proposes a parameter learning-based method to tackle LTR tasks. [52] proposes a mixture feature transformation mechanism which could automatically derive a mixture of basic feature transformation functions to optimize ranking performance. Nevertheless, under appropriate settings, over-parameterization could gain better performance on test data in the newly-fashioned “interpolation regime” with the *double descent curve*. There are several over-parameterization methods [33], [34] which have superb performance. Random Fourier Feature [53] adopts a kernel technique to generate features for most inner product-based models, which has gained great improvements. COLTR follows this line of research and is the first to leverage RFF-based over-parameterization to improve LTR models.

### 2.3 Semi-supervised Learning

Nowadays, semi-supervised learning methods have been widely adopted in machine learning tasks, such as classification, regression, etc. Two effective categories of semi-supervised learning methods are self-training [14], [12] and co-training [54]. For self-training, the basic idea is to generate pseudo-labels for unlabeled data and improve the performance with pseudo-labeled data [55]. The learning process of semi-supervised learning could follow the four-step strategy: (1) training a model with labeled data; (2) using the trained model to generate pseudo-labels; (3) adding the pseudo-labeled data into the label data and training another model with the combined data; (4) retraining the former model with labeled data. [56], [57], [58] present the effectiveness of co-training and gain significant improvements. Moreover, [51] proposes a semi-supervised learning method for LTR tasks with preference regularization. Therefore, it is illustrated that the co-training method is useful for web-scale search. *In this work, COLTR adopts co-training [27], [28], [29] for semi-supervised LTR, where listwise and pointwise models teach each other based on pseudo-labels via predictions.*

## 3 COLTR DESIGN: PRELIMINARIES

In this section, we introduce preliminary works for our proposed model. We first detail feature construction and retrieval and ranking candidates for COLTR. Then, we formulate semi-supervised LTR problems for our proposed model. Table 1 lists the notations and their definitions.

### 3.1 Feature Construction & Retrieval and Ranking Candidates for COLTR

Given the massive webpages archived and indexed (please see also in Section 5 for the deployment of COLTR, where storage & indexing would be introduced), compared with traditional approaches, such as text matching, COLTR leverages pre-trained language models [59] based semantic retrieval algorithms to conduct an effective and efficient online webpages retrieval with given queries, and provide features extracted from webpages/queries for LTR.

TABLE 1: List of Notations

Notation	Definition
[CLS]	pseudo token for the subsequent matching
$\{T_1, \dots, T_N\}$	tokenized sequence of the raw query
$\{T'_1, \dots, T'_N\}$	tokenized sequence of the raw webpage
$\{F_1, \dots, F_N\}$	encoded embedding of the query
$\{F'_1, \dots, F'_N\}$	encoded embedding of the webpage
$\mathcal{Q}$	set of search query
$\mathcal{D}$	set of all archived webpages
$q_i$	the $i^{\text{th}}$ query in $\mathcal{Q}$
$d_i$	the $i^{\text{th}}$ webpage in $\mathcal{D}$
$D_i$	the set of relevant webpages for $q_i$
$\mathbf{y}_i$	set of ranking scores for $q_i$
$y_j^i$	ranking score of the $j^{\text{th}}$ webpage for $q_i$
$d_j^i$	the $j^{\text{th}}$ webpage of $D_i$
$\mathcal{T}$	set of query-webpage pairs with ranking scores
$\mathcal{T}'$	set of unlabeled query-webpage pairs
$\mathbf{x}_{i,j}$	feature vector of the query-webpage pair $(q_i, d_j^i)$
$m$	original dimensions of feature vectors
$\mathbf{z}_{i,j}$	transformed feature vector
$N$	transformed dimensions of feature vectors
$\mathcal{Z}^L$	set of transformed labeled query-webpage pairs
$\mathcal{Z}^U$	set of transformed unlabeled query-webpage pairs
$\mathcal{Z}^P$	set of pseudo-labeled data
$\mathcal{Z}^C$	set of combined data
$\text{LTR}^{Po}$	pointwise-based LTR model
$\text{LTR}^{Li}$	listwise-based LTR model
$C$	the number of rounds for co-training

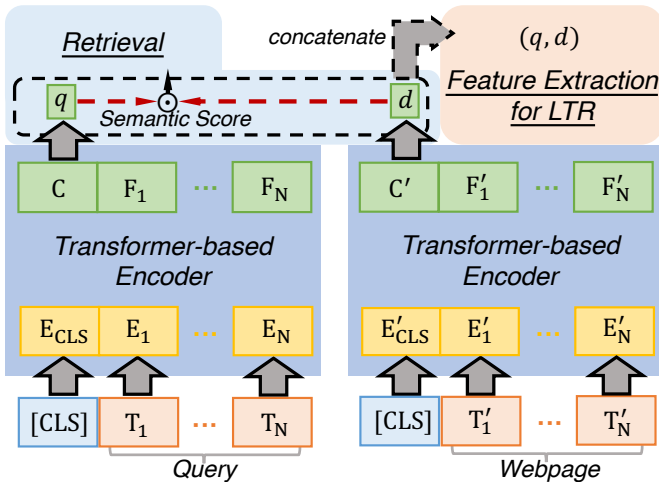


Fig. 1: The Process of Feature Extraction and Retrieval with ERNIE-based Semantic Retrieval Model. The Transformer-based encoders take the tokenized sequence of the query and webpage as the input and output encoded query and webpage embeddings, which are used to do semantic retrieval and conduct the query-webpage pair feature for LTR.

Specifically, we adopt the ERNIE-based [60] semantic retrieval model to enhance the conventional retrieval approach [61]. For conventional retrieval, it performs numerous operations on query texts, such as word segmentation and stop-word filtering, utilizes term indexes to accomplish keyword matching, and gets a set of relevant webpages. Meanwhile, the retrieval method based on pre-trained models first leverages an ERNIE transformer module to obtain the embeddings of queries and webpages. More specifically, the transformer encoder [62] first takes the tokenized sequence of the raw query or webpage, i.e.,  $\{[CLS], T_1, \dots, T_N\}$  or  $\{[CLS], T'_1, \dots, T'_N\}$ , as the input, where the

pseudo token [CLS] aggregates data in the encoder for the subsequent matching, and outputs an encoded embedding of the query or webpage, i.e.,  $\{C, F_1, \dots, F_N\}$  or  $\{C, F'_1, \dots, F'_N\}$  in Figure 1. Note that, the retrieval model together with the ERNIE module is end-to-end trainable given a bucket of feedback, such as click-throughs, dwell time, and the raw text-to-text matching results by conventional retrievals, as supervision signals [26]

Given a query and a webpage, the ERNIE-based retrieval estimates a semantic score between the query and webpage through computing the cosine similarity or inner-product between their embeddings. Then, under a query, COLTR tops the webpages with highest semantic scores as the results of retrieval and also the **candidates for webpages ranking**. Then COLTR forwards embeddings of the query and the top-retrieved webpages (as well as their semantic scores and other statistical features, e.g., click-through of webpages and etc.) as the feature inputs for LTR.

### 3.2 Semi-supervised LTR Problems for COLTR

In this section, we formalize the LTR settings and propose our notations. Given a set of search queries  $\mathcal{Q} = \{q_1, q_2, \dots\}$  and all archived webpages  $\mathcal{D} = \{d_1, d_2, \dots\}$ , for each query  $q_i \in \mathcal{Q}$ , the search engine could retrieve a set of relevant webpages denoted as  $D_i = \{d_1^i, d_2^i, \dots\} \subset \mathcal{D}$ . Through annotating, there also might exist a set of relevance labels  $\mathbf{y}_i = \{y_1^i, y_2^i, \dots\}$  for  $q_i$ , which characterizes the relevance of each document  $d_j^i \in D_i$  to the search query  $q_i$ . In our work, we follow the settings in [22], [4] and scale the relevance label from 0 to 4 to represent levels of relevance (i.e., **{bad, fair, good, excellent, perfect}** the bigger the more relevant).

#### 3.2.1 Learning Objective of LTR

We denote a set of query-webpage pairs with ranking score annotations as a set of triples such as  $\mathcal{T} = \{(q_i, D_i, \mathbf{y}_i)\}_{i=1}^{|\mathcal{T}|}$ . We aim to gain an LTR scoring function  $f: \mathcal{Q} \times \mathcal{D} \rightarrow [0, 4]$ . Therefore, the goal is recast to learn a scoring function  $f$  which minimizes the loss as:

$$\mathcal{L}(f) = \frac{1}{|\mathcal{T}|} \sum_{i=1}^{|\mathcal{T}|} \ell(\mathbf{y}_i, F), \quad (1)$$

where  $F = \{f(q^i, d_j^i)\}_{j=1}^{|D_i|}$  is the set of webpage ranking scores, and  $\mathbf{y}_i$  is the set of scale with  $y_j^i$  representing the relevance label corresponds to  $d_j^i$ .  $\ell$  represents the average loss function of the ranking predictions of all archived webpages  $D_i$  of query  $q_i$  against the ground truth  $\mathbf{y}^i$ . Three types of loss functions are defined as follows.

- 1) The **pointwise** loss converts LTR into a regression task with a L2-loss such that  $\ell = \frac{1}{|D_i|} \sum_{j=1}^{|D_i|} |f(q^i, d_j^i) - y_j^i|^2$  for webpage  $d_j^i$  of query  $q_i$ .
- 2) The **pairwise** loss is based on the pairwise order of any two webpages  $\{d_j^i, d_k^i\} \subset D_i$ . Here, we follow[63], such that:

$$\ell = \frac{1}{|D_i|} \sum_{j=1}^{|D_i|} \log \left\{ 1 + e^{(y_j - y_k)} \right\} \cdot |\Delta Z_{jk}|, \quad (2)$$

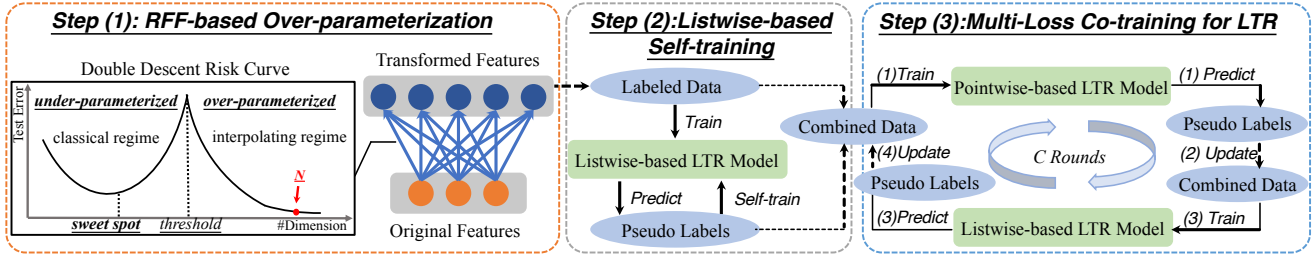


Fig. 2: The Pipeline of COLTR consisting of three steps: (1) *RFF-based Over-parameterization*, (2) *Listwise-based Self-training*, and (3) *Multi-Loss Co-training for Semi-supervised LTR*.

where  $\Delta Z_{jk}$  reflects the effects [40] by swapping the positions of two webpages.

- 3) The **listwise** loss measures, for each query  $q_i$ , the divergence between the probability distributions of the predicted ranking scores for every webpage in  $D_i$  and the ground truth  $y_i$ . The divergence is estimated using cross-entropy, while the probability distribution of ranking scores is estimated using softmax-like normalization [64].

### 3.2.2 Semi-supervised LTR

As annotators for web search can only label a small number of query-webpage pairs while annotating more query-webpage pairs might be expensive and time-consuming, the core problem of LTR is thus to incorporate some unlabeled query-webpage pairs, i.e.,  $\mathcal{T}' = \{(q'_1, D'_1), (q'_2, D'_2), \dots\} \subset \mathcal{Q} \times 2^{\mathcal{D}}$  and  $|\mathcal{T}'| \gg |\mathcal{T}|$ , to improve training. To this end, the *Semi-supervised Learner* plays to train LTR models using both labeled and unlabeled query-webpage pairs. In our research, we propose COLTR to enable semi-supervised LTR with pseudo-labels via predictions.

## 4 COLTR DESIGN: ALGORITHMS

In this section, we introduce the algorithm design of our proposed model COLTR. As illustrated in Figure 2, COLTR consists of three steps: (1) *RFF-based over-parameterization*, (2) *Listwise-based Self-training*, and (3) *Multi-Loss Co-training for LTR*. Specifically, we first introduce *RFF-based over-parameterization* in Section 4.1, and then describe *Listwise-based Self-training* and *Multi-Loss Co-training for LTR* together in Section 4.2.

### 4.1 RFF-based Over-parameterization

Given the overall set of queries  $\mathcal{Q}$  and the set of all webpages  $\mathcal{D}$ , COLTR first obtains every possible query-webpage pair from both datasets, denoted as  $(q_i, d_i^j)$  for  $\forall q_i \in \mathcal{Q}$  and  $\forall d_i^j \in D_i \subset \mathcal{D}$ , i.e., the  $j^{\text{th}}$  webpage retrieved for the  $i^{\text{th}}$  query. For each query-webpage pair  $(q_i, d_i^j)$ , COLTR further extracts an  $m$ -dimensional feature vector  $x_{i,j}$  representing the features of the  $j^{\text{th}}$  webpage under the  $i^{\text{th}}$  query, using the pre-trained language models [4], [26]. For more details about feature extraction, please refer to Section 3.1 and Figure 1. The goal to incorporate RFF-based over-parameterization is to further enhance the learned representations from the language model, in order to train LTR models in the interpolating regime [65].

Given  $x_{i,j} \in \mathcal{R}^m$ , COLTR further maps the feature vector into an  $N$ -dimensional vector denoted  $z_{i,j} = \mathbf{z}(x_{i,j})$

### Algorithm 1: Random Fourier Feature with Gaussian Kernel

**Output:** the function  $\mathbf{z}(x) : \mathcal{R}^m \rightarrow \mathcal{R}^N$ .

**begin**

Define Gaussian Distribution:

$$p(\omega) = (2\pi)^{-\frac{N}{2}} e^{-\frac{\|\omega\|_2^2}{2}};$$

Draw  $N$  i.i.d samples  $\omega^1, \dots, \omega^N \in \mathcal{R}^m \sim p(\omega)$ ;

Draw  $N$  i.i.d samples  $b^1, \dots, b^N \in \mathcal{R} \sim \mathcal{N}(0, 1)$ ;

Conduct  $W \in \mathcal{R}^{N \times m} = \{\omega^1, \dots, \omega^N\}$  and

$B \in \mathcal{R}^N = \{b^1, \dots, b^N\}$ ;

Compute  $\mathbf{z}(x) = \sqrt{\frac{2}{N}} [\cos(W^T x + B)]$ ;

**return**  $\mathbf{z}(x)$ ;

**end**

using the feature transformation function  $\mathbf{z}(x)$ . In this work, we use the transformation based on Random Fourier Features [53] to implement  $\mathbf{z}(x)$ , which is defined in Algorithm 1. Please be advised that the use of  $\mathbf{z}(x)$  can map original features of LTR into a feature space of higher dimensions when  $N \gg m$ . With the increasing number of dimensions  $N$ , the LTR model is being over-parameterized with more input features and would incorporate feature-wise “double descent” phenomenon of generalization errors in prediction [33], [34]. Through cross-validation on labeled set  $\mathcal{T}$ , COLTR determines the optimal setting of  $N$  to achieve the best generalization performance. Hence, with  $z_{i,j}$  for every query-webpage pair, the over-parameterized LTR model is expected to work in the interpolating regime [34] with superb generalization performance.

In this way, COLTR transforms query-webpage pairs in labeled and unlabeled datasets  $\mathcal{T}$  and  $\mathcal{T}'$  into two sets of labeled and unlabeled feature vectors  $\mathcal{Z}^L = \{(z_{i,j}, y_j^i) | \forall (q_i, D_i, \mathbf{y}) \in \mathcal{T} \text{ and } \forall d_i^j \in D_i\}$  and  $\mathcal{Z}^U = \{z_{i,j} | \forall (q_i, D_i) \in \mathcal{T}'\}$ , respectively. Note that, compared to raw features directly extracted from language models, RFF-based over-parameterization would significantly overload the LTR regressor with a larger number of features. However, compared to the end-to-end training based on the language model and LTR regressor, RFF-based over-parameterization is with low cost while it still performs LTR in the interpolating regime.

### 4.2 Listwise-based Self-training and Multi-Loss Co-training for Learning to Rank

Given the labeled and unlabeled sets of feature vectors  $\mathcal{Z}^L$  and  $\mathcal{Z}^U$ , COLTR further takes the next two steps in the pipeline to accomplish semi-supervised LTR. Algorithm 2

---

## Algorithm 2: Listwise-based Self-training and Multi-Loss Co-training for LTR

---

**Input:** labeled data  $\mathcal{Z}^L$ , unlabeled data  $\mathcal{Z}^U$ , the number of rounds for co-training  $C$ .  
**Output:** Trained pointwise model  $\text{LTR}^{Po}$

```

begin
  /* Listwise-based Self-training */
  Train listwise model  $\text{LTR}^{Li}$  on  $\mathcal{Z}^L$ ;
   $\mathcal{Z}^P \leftarrow$  pseudo-labels  $\mathcal{Z}^U$  using predictions of
   $\text{LTR}^{Li}$ ;
   $\mathcal{Z}^C \leftarrow \mathcal{Z}^P \cup \mathcal{Z}^L$ ;
  /* Multi-Loss Co-training for LTR */
  for  $i = 1$  to  $C$  do
    Train pointwise model  $\text{LTR}^{Po}$  on  $\mathcal{Z}^C$ ;
     $\mathcal{Z}^P \leftarrow$  pseudo-labels  $\mathcal{Z}^U$  using predictions of
     $\text{LTR}^{Po}$ ;
     $\mathcal{Z}^C \leftarrow \mathcal{Z}^P \cup \mathcal{Z}^L$ ;
    Train listwise model  $\text{LTR}^{Li}$  on  $\mathcal{Z}^C$ ;
     $\mathcal{Z}^P \leftarrow$  pseudo-labels  $\mathcal{Z}^U$  using predictions of
     $\text{LTR}^{Li}$ ;
     $\mathcal{Z}^C \leftarrow \mathcal{Z}^P \cup \mathcal{Z}^L$ ;
  end
  return  $\text{LTR}^{Po}$ ;
end

```

---

lists the pseudo-codes of the two steps of *Listwise-based Self-training* and *Multi-Loss Co-training for LTR*, respectively.

### 4.2.1 Listwise-based Self-training

First of all, **COLTR** trains a listwise LTR model  $\text{LTR}^{Li}$  using both  $\mathcal{Z}^L$  and  $\mathcal{Z}^U$  through self-training, where  $\text{LTR}^{Li}$  is first trained using  $\mathcal{Z}^L$  through supervised learning. Then,  $\text{LTR}^{Li}$  predicts the ranking score for each feature vector in  $\mathcal{Z}^U$  and *pseudo-labels* the feature vector with the prediction result. Later, **COLTR** combines  $\mathcal{Z}^L$  with pseudo-labeled data  $\mathcal{Z}^P$  and retrains  $\text{LTR}^{Li}$  using the *combined data*  $\mathcal{Z}^C$ .

### 4.2.2 Multi-Loss Co-training for LTR

Given the *combined data*,  $\mathcal{Z}^L$  and  $\mathcal{Z}^U$ , **COLTR** (1) trains a pointwise model  $\text{LTR}^{Po}$  using  $\mathcal{Z}^C$  and predicts pseudo-labels for every feature vector in  $\mathcal{Z}^U$  using trained  $\text{LTR}^{Po}$ . Later, **COLTR** (2) updates  $\mathcal{Z}^P$  with the prediction results of  $\text{LTR}^{Po}$  and combines  $\mathcal{Z}^L$  with  $\mathcal{Z}^P$  to obtain  $\mathcal{Z}^C$ . **COLTR** further (3) retrains  $\text{LTR}^{Li}$  using the  $\mathcal{Z}^C$  and predicts ranking scores for each feature vector in  $\mathcal{Z}^U$  using trained  $\text{LTR}^{Li}$ . Finally, **COLTR** (4) updates  $\mathcal{Z}^P$  with the prediction results of  $\text{LTR}^{Li}$  and combines  $\mathcal{Z}^L$  with  $\mathcal{Z}^P$  to obtain  $\mathcal{Z}^C$ . **COLTR** repeats above (1)–(4) steps with  $C$  rounds and uses the pointwise model  $\text{LTR}^{Po}$  to serve online ranking.

## 4.3 Discussion on COLTR Algorithms

In this section, we discuss the research work of **COLTR** from the following perspectives.

### 4.3.1 Features from the Language Models

In this work, we design LTR algorithms and **COLTR** on top of a pre-trained language models-based retrieval method [26], which supplies the candidates for webpage ranking and the features of LTR. The retrieval models

together with ERNIE transformers are end-to-end trained incorporating both self-supervision of language models and some supervision signals, including text-to-text matching, click-throughs and dwell time. **COLTR** however doesn't include LTR models in the end-to-end training with the language models, primarily due to the cost reason. In fact, the requests to re-training LTR models would be highly frequent due to the fast shift the internet interests, however, the language models would be updated infrequently as the essentials of languages would not change rapidly. In this way, **COLTR** can train and re-train LTR models in a low cost while enjoying the features extracted from language models.

### 4.3.2 LTR with Over-parameterization

We incorporate *over-parameterization* in LTR to improve the generalization and its performance of generalization. In recent years, *Belkin et al* [33] together with other researchers have found a new phenomenon in machine learning, namely *double descent*, where after a period of increasing complexity, the performance of a model doesn't degrade and starts to improve again, leading to a second descent in the generalization error curve. The classic statistical learning theory, such as VC dimensions [66], models with lower complexity tend to perform better (in the classic regime) and achieves the right balance between bias and variance at a point (i.e., sweet spot), where the model is neither underfitting nor overfitting the data. However, double descent suggests that there can be significantly better generalization performance for a model with higher complexity than what VC dimensions predict (in the so-called interpolation regime) [65], as the model can "perfectly" fits the data (i.e., interpolation) while avoiding over-fitting [34]. Therefore, double descent challenges the traditional view that models should be kept as simple as possible to avoid over-fitting and suggests that more complex models may be desirable in certain cases (such as deep learning).

The complexity here can include the number of features, the number of parameters, the number of data points, or the number of iterations used to train a model [67], [68], [69]. For **COLTR**, we adopt *Random Fourier Features* (RFF) [53] to extend the number of features and train the gradient boosting regression trees with transformed data — the use of RFF to over-parameterize statistical learners has been widely studied or practiced in double descent works [33], [67], [70], [71], [72]. In terms of theory to support **COLTR** design, we recommend readers to refer to the analytical studies [70], [71], where the generalization error bounds of RFF in interpolation regime have been studied. In this section, we do not repeat the theoretical analysis here as they are just off-the-shelf.

Of-course, in past research work, Random Fourier Features have been also widely used to facilitate (Gaussian) Kernel machines [53], [73], [74] for potential generalization performance improvement. However, our RFF-based over-parameterization is quite different from previous works, as the traditional approach adopts RFF to reduce the dimensions of features while **COLTR** adopts RFF as a way to enhance learned representations by increasing the number of dimensions. Recent studies [75] show that RFF with sufficient feature extension (over-parameterization) can bring significant performance improvement.




Fig. 3: The Overview of Baidu Search with COLTR Deployed.

### 4.3.3 Semi-supervision with Multi-Loss Co-training

COLTR leverages multiple LTR regressors based on different losses (e.g., pointwise, pairwise, and listwise losses) to co-train each other via pseudo labels. The goal of incorporating different losses is to introduce diversity into semi-supervision signals. More specifically, COLTR pipes the pointwise-based LTR regressor and listwise-based LTR in the loop of co-training, as the diversity between two regressors is large — one predicts the score of ranking while the other considers the position in a list. Our latter offline and online experiments demonstrate the advantage of such configuration on real-world data, compared to other combinations. Note that, co-training by diversity between learners is a common practice in semi-supervised learning and has been studied extensively from both theoretical and practical perspectives [58], [27], [76].

## 5 COLTR DEPLOYMENT

In this section, we introduce the deployment settings of COLTR at Baidu Search. As shown in Figure 3, the industrial search engine is essential with three stages as follows: (1) *Data collection*, (2) *Storage & Indexing* and (3) *Retrieval & ranking*.

### 5.1 Data Collection

As a fully-automated search engine, Baidu search uses software that explores the web on a regular basis to find sites to be added to the database. In fact, the vast majority of sites listed in the results are not manually submitted for inclusion but are found and added automatically when the web crawler crawls the web. On the contrary, they are discovered and uploaded automatically as the web crawler browses the internet. Given the massive webpages on the web, Baidu search engine adopts a high-performance crawler, namely Web Crawler, to fetch and download webpages from the web. Specifically, the Web Crawler screens a list of links (*i.e.*, URLs) for new webpages and updated ones, then stores the valid links (or URLs) with desired contents in a large downloading list. Later, the Web Crawler starts downloading the webpages on the list, upon the real-time web traffics of the Baidu search engine. Note that once a new or updated webpage is fetched, the Web Crawler would first parse the contents, find all possible links, and add them into the list for potential screening.

### 5.2 Storage & Indexing

Given the massive webpages downloaded from the web, Baidu search stores these contents in distributed archival storage systems with Fatman [77] and builds efficient indices for high-performance search based on DirectLoad [78]. While Fatman [77] could significantly reduce the costs of storage by using the elastic resources, such as underutilized servers and temporally spare storage, across multiple regional data centers of Baidu, DirectLoad [78] balances loads of indexing over these data centers with superb I/O efficiency using novel key-value operations and in-memory computation.

### 5.3 Retrieval & Ranking

For all webpages in the storage, as was mentioned in Section 3.1, Baidu search leverages an ERNIE-based encoder to compute the embeddings and conducts an embedding database, which will be used in the retrieval stage. Please refer to Section 3.1 for details for webpages retrieval.

In terms of webpages ranking in online settings, as was mentioned in Section 4, COLTR totally obtains three models  $LTR^{Po}$ ,  $LTR^{Li}$ , and  $LTR^{Pa}$  through  $C$  rounds of co-training via pointwise, listwise and pairwise losses, respectively. COLTR adopts  $LTR^{Po}$  to serve the online ranking tasks. More specifically, after the training procedure of LTR, COLTR restored the RFF-transformation of LTR features for either webpages or queries as the immediate results of online ranking. For online inference, given a query and a webpage for online ranking, COLTR pickups their RFF-transformed representations and passes them to the GBM for inference in a fast manner at Baidu Search.

## 6 EXPERIMENTS

To demonstrate the effectiveness of our proposed model, we present extensive experiments compared with a large number of baseline methods. Firstly, we introduce the experimental details in terms of the dataset, evaluation methodology, competitor system, and experimental settings. Then, we introduce the results of offline experiments. Finally, the online A/B Test performance shows the effectiveness of COLTR at Baidu Search.

### 6.1 Experimental Details

#### 6.1.1 Dataset

We evaluate the proposed model and baseline models on the dataset collected from Baidu search engine system. Due to privacy concerns, it should be noted that none of the data contains any user-related information. Specifically, the dataset consists of 15,000 queries and over 770,000 query-webpage pairs. The dataset is split into training set (12,000 queries), validation set (1,000 queries) and test set (2,000 queries), which contain 616,314 query-webpage pairs, 51,360 query-webpage pairs and 103,872 query-webpage pairs, respectively. For the semi-supervised learning experiments, we randomly select four ratios of labeled data from training set, where we utilize  $\alpha = \{0.05, 0.1, 0.15, 0.2\}$  to represent the four ratios, respectively. Table 2 shows the statistics of the dataset.

TABLE 2: Statistics of the dataset.

Dataset	#Query	#Query-webpage pairs
Training Set	12,000	616,314
Validation Set	1,000	51,360
Test Set	2,000	103,872

### 6.1.2 Evaluation Methodology

To evaluate the performance of our proposed method, we use Normalized Discounted Cumulative Gain (NDCG) [40], which has been widely adopted to evaluate the relevance in the context of ad-hoc search engine. For a query and its relevant webpages, the LTR model usually predicts a score for each webpage and generates a ranking list by sorting scores in descending order. *From the perspective of research and business to evaluate the models' performance, we consider the NDCG of the top 10 and 4 ranking results, i.e.,  $NDCG_{10}$  and  $NDCG_4$ .*

### 6.1.3 Competitor Systems and Baselines

For all experiments, the baseline model is a *pointwise-based self-trained LTR model without RFF-based over-parameterization*. Specifically, the baseline model adopts a self-trained LightGBM [18]-based LTR model with an L2 loss function, which has been deployed at Baidu search. Moreover, in order to demonstrate the effectiveness of COLTR sufficiently, we choose seven semi-supervised LTR models as the comparative models. Specifically, we utilize "Po", "Pa" and "Li" to represent RFF-overparameterized *self-trained* LTR models of pointwise, pairwise and listwise, respectively. We use "Po-Pa", "Po-Li", "Pa-Po", "Pa-Li", and "Li-Pa" to represent the RFF-overparameterized *multi-loss co-training* LTR models under pointwise-to-pairwise, pointwise-to-listwise, pairwise-to-pointwise, pairwise-to-listwise, and listwise-to-pairwise settings, where the terminology is based on the order of LTR models used for co-training, such that COLTR is exactly in the listwise-to-pointwise ("Li-Po") setting. We choose the pointwise-based self-trained LightGBM-based LTR model without *RFF-based over-parameterization* as the *base* model. In this work, due to the restriction of business information disclosures, we only report  $\Delta NDCG$  to measure the difference between our proposed model and the *base* model.

### 6.1.4 Experimental Settings

In this work, all the offline experiments are implemented on PaddlePaddle<sup>2</sup> cloud platform with 64G Memory, 4 NVIDIA Tesla V100 GPU, and 12T Disk. The online experiments are deployed in Baidu search engine system. More online A/B Test setting details are introduced in Section 6.3.1. We choose LightGBM, which is the most popular tree-based ranker, as the base ranking model with the number of trees as 2000 and the learning rate as 0.01. The experiments consist of nine self-training and co-training models under four ratios of labeled data for ten rounds. Besides, for the *RFF-based over-parameterization* experiments, we set the ratio of the number of transformed dimension  $N$  and the number of original dimensions  $m$  as  $N/m$ .

2. <https://www.paddlepaddle.org.cn/>

## 6.2 Offline Experimental Results

To comprehensively evaluate our proposed method, we conduct experiments to answer the following questions:

**RQ1** How does COLTR perform compared with the baseline for LTR tasks?

**RQ2** Which number of the transformed dimension can make LTR models gain the best performance under different ratios?

**RQ3** Is the *RFF based over-parameterization* in COLTR necessary for improving performance?

**RQ4** How does the number of rounds impact the performance of COLTR?

### 6.2.1 Comparative Results: RQ1

In Figure 4, we report the offline performance of COLTR compared with other baselines under four different ratios of labeled data on  $\Delta NDCG_{10}$  and  $\Delta NDCG_4$ . For each semi-supervised learning model, we choose the model with the best performance of all validation rounds for testing. Intuitively, we could see that COLTR gains the best performance compared with other baselines under four ratios of labeled data on both two metrics. Specifically, COLTR achieves 4.92% improvement on  $\Delta NDCG_4$  under 5% ratio of labeled data. COLTR incorporates the diversity of prediction results of listwise model and pointwise model in a loop of multiple rounds. As the stronger learner, listwise model predicts more accurate pseudo-labels for pointwise model. Then the weaker model, pointwise model, generates relatively inaccurate but diverse pseudo-labels to train the stronger model. In such *Multi-Loss Co-training* mechanism, COLTR gains significant performance. Moreover, there are two findings in the comparative results. First, when  $\alpha = 5\%$ , COLTR significantly outperforms other self-training models and co-training models compared with other three ratios of labeled data on  $\Delta NDCG_4$ . Next, although co-training models can not obtain the performance like COLTR, some models, such as "Pa-Po", "Pa-Li", "Li-Pa", also outperform three self-training models when  $\alpha = \{0.1, 0.15, 0.2\}$  on  $\Delta NDCG_{10}$ .

### 6.2.2 Ablation Study: RQ2

In this study, we conduct a series of experiments to investigate the proper value of transformed dimension under four kinds of ratios of labeled data for LTR tasks. In order to show the results clearly, we choose the performance on  $\Delta NDCG_{10}$  instead of "test error" to present the curves. Intuitively, Figure 5 represents the margin curves for the pointwise model under four different ratios of labeled data on validation set. In Figure 5 (a), we present that the pointwise model gains the best performance when  $N/m = 17$  under 5% labeled training data. As depicted in the figure, at the beginning of the curve, the value of  $\Delta NDCG_{10}$  first rises and then falls in the "classical regime". Once the value of  $N/m$  exceeds the threshold, the pointwise model performs well again and gains the best performance at  $N/m = 17$  in the "interpolating regime". For different ratios of labeled data, the value of  $N/m$  is various in terms of the number of labeled samples. Figure 5 (b) shows the most proper value of  $N/m$  is 21 under 10% of labeled data. Similarly, when  $\alpha = 0.15$ , the chosen value of  $N/m$  is






Fig. 4: Offline comparative results ( $\Delta NDCG_{10}$  and  $\Delta NDCG_4$ ) of **COLTR** and baselines under various ratios of labeled data.




Fig. 5: Ablation studies of *RFF-based Over-parameterization* to choose the proper transformed dimensions under various ratios of labeled data on the validation set. We choose the values of  $N/m$  marked by dashed lines.

25. For 20% proportion of label data, the chosen value of  $N/m$  is 30. We can see that transforming the original query-webpage vector into the proper dimension could improve the performance of downstream LTR models in the newly-fashioned “interpolating regime”.

### 6.2.3 Ablation Results: RQ3

In order to demonstrate the effectiveness of *RFF-based Over-parameterization* part of **COLTR**, we conduct the ablation study of *RFF-based Over-parameterization* for **COLTR** and baselines under various ratios of labeled data. As we can see, all the semi-supervised learning models with *RFF-based Over-parameterization* could obtain better performance compared with semi-supervised learning models without *RFF-based Over-parameterization*. As depicted in Figure 6 (b), *RFF-based Over-parameterization* achieves the improvement with 3.30% for **COLTR** in average under 5% proportion of labeled data on  $\Delta NDCG_4$ , which is the largest improvement of **COLTR**. Besides, the pointwise-based self-training model with *RFF-based over-parameterization* obtains the improvement with 3.38% on average, which is the largest improvement among all experiments. Semi-supervised learning LTR model without *RFF-based Over-parameterization* takes a small number of hand-crafted features and is based on tree-based models and their derivatives to pursue high throughput under concurrency poorly. Obviously, *RFF-based Over-parameterization* could tackle that issue well.

### 6.2.4 Parameter Sensitivity: RQ4

In this section, we conduct a series of experiments to study the performance variation of **COLTR** with respect to the number of rounds for co-training. In Figure 7, we report the

studies of **COLTR** under various ratios of labeled data in ten co-training rounds on the validation set. The results show that **COLTR** gains the best performance at the 5th round in all experiments. For each round, the listwise model is trained on the combined data and generates pseudo-labels. Next, the pseudo-labeled data is combined with the labeled data. The pointwise model is trained on the combined data and generates the results on  $\Delta NDCG$ . As we can see in Figure 7 (a), with the number of rounds increasing, **COLTR** gains a rising performance and obtains the best performance at the 5th round. Next, the performance starts to decrease and is in a state of shock. Finally, the best results can be obtained at the 5th round. The other experiments have similar phenomena like Figure 7 (a). Therefore, we choose the trained **COLTR** at the 5th round for online experiments.

## 6.3 Online Experimental Results

To investigate the impact of **COLTR** at Baidu Search, we deploy the new system and conduct a series of online A/B Test with real-world web traffics compared with the base model in Baidu Search.

### 6.3.1 A/B Test Setups

In the online A/B Test, we conduct the experiment that compares the new ranking system, which deploys **COLTR**, with the old system for 7 days. For each day, we first remove pornographic and legally prohibited webpages. Then, we hire six common annotators to annotate the relevant score for each chosen query-webpage pair. Next, our professional annotators evaluate the quality of the annotations and guarantee that the accuracy is higher than 85%. Eventually, we leverage the weighted average of the annotations as the

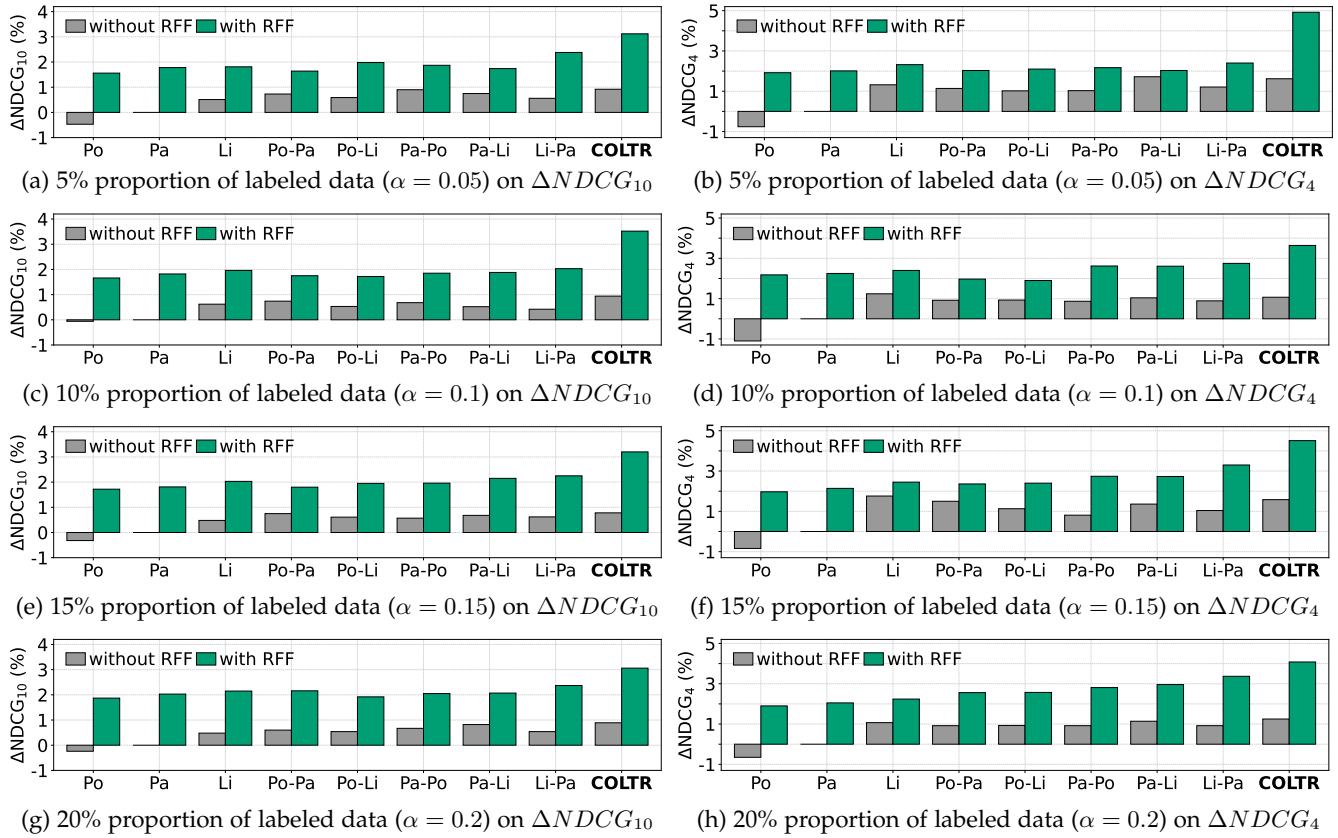


Fig. 6: Ablation studies of *RFF-based Over-parameterization* for **COLTR** and baselines under different ratios of labeled data.

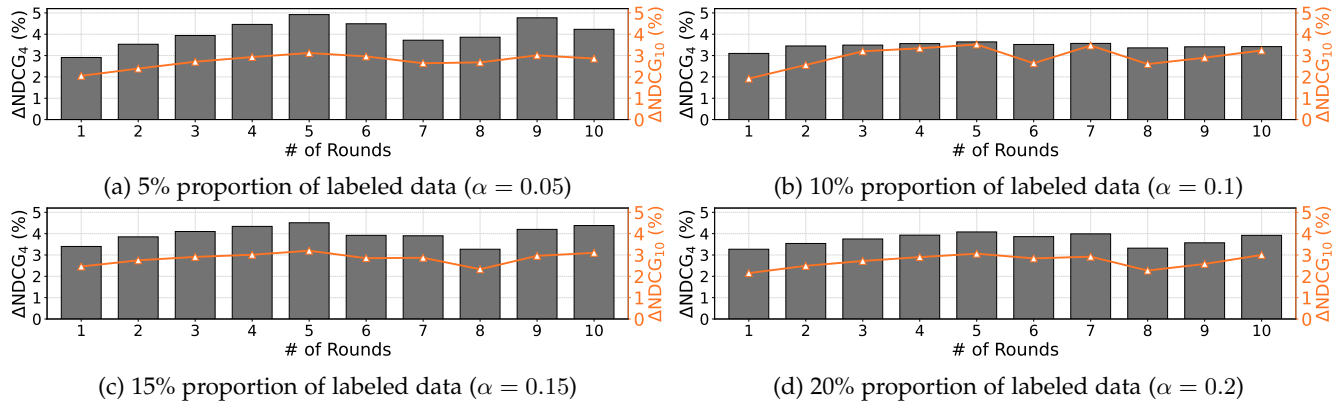


Fig. 7: Performance ( $\Delta NDCG_4$  and  $\Delta NDCG_{10}$ ) of **COLTR** under four ratios of labeled data in ten rounds on the validation set.

relevant score to train our proposed model. According to the offline experimental results, we choose the trained **COLTR** and other baseline models under four various ratios of labeled data in the best performance round. We conduct the online experiments with 0.6% real-world web traffics of Baidu search and concentrate on metrics that have a direct impact on user experience. From the perspective of business, we consider the NDCG of the top 4 ranking results and calculate  $\Delta NDCG_4$  between the chosen model and the online base model.

### 6.3.2 Online Performance

Figure 8 illustrates the comparison of **COLTR** with the baselines on  $\Delta NDCG_4$ . Firstly, **COLTR** could boost the

performance compared with the online base system in all days, which demonstrates **COLTR** is practical for improving the performance of Baidu search engine. Furthermore, we can find that **COLTR** achieves significant improvements in Baidu search engine. Specifically, we observe that **COLTR** trained on 5% proportion of labeled data outperforms the online base model by a large margin on  $\Delta NDCG_4$  with 0.92% relative improvement. The largest improvements of trained **COLTR** on 10%, 15% and 20% proportion of labeled data are 0.76%, 0.72% and 0.65% on  $\Delta NDCG_4$ , respectively. These significant improvements reveal the effectiveness of **COLTR**. Finally, we also compare **COLTR** with other semi-supervised learning models used in our offline experiments. We notice that **COLTR** outperforms several semi-supervised

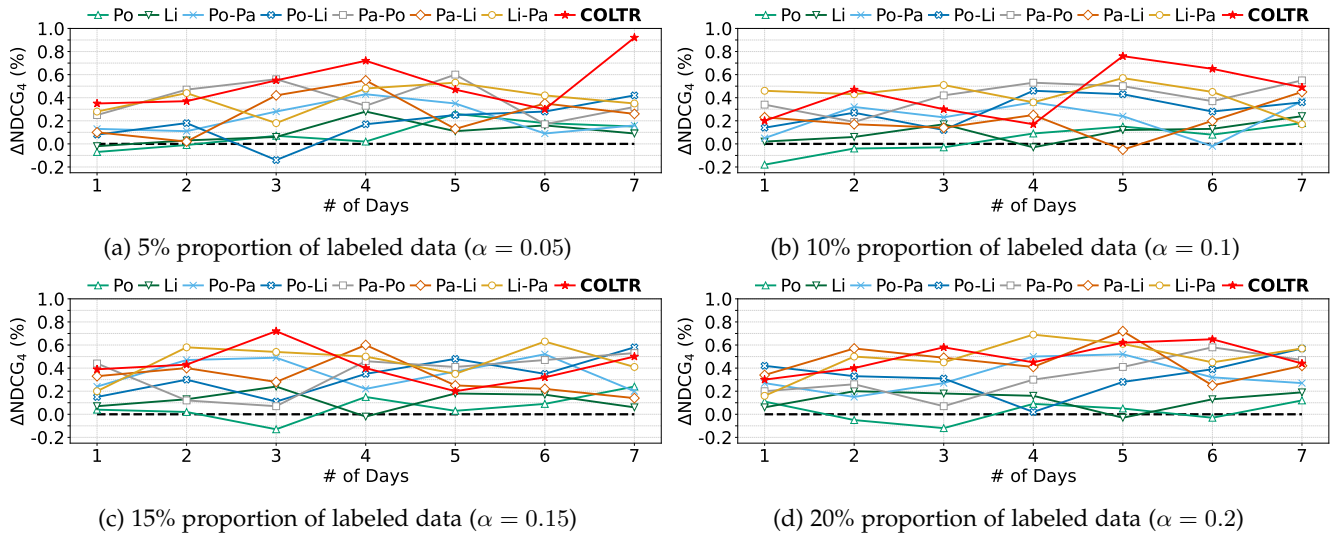


Fig. 8: Online comparative performance ( $\Delta NDCG_4$ ) of COLTR and baselines for 7 days ( $t$ -test with  $p < 0.05$  over the baseline). COLTR could boost the performance compared with the online base system in all days, which demonstrates COLTR is practical for improving the performance of Baidu search engine.

learning models in all days, which proves that COLTR is more useful and sound to improve the accuracy of real-world online search engine. Basically, the online performance is consistent with our offline experiment result.

#### 6.4 Discussion on Experiments

In our experiments, we adopt LightGBM to generate pseudo labels and predict ranking scores. Actually, the main contribution of this work is a semi-supervised LTR framework leveraging features extracted from large language models, generating pseudo labels by multi-loss co-training, and enhancing rank regressor via Over-parameterization. Many detailed configurations of COLTR are indeed interchangeable, such as the adoption of LightGBM to predict ranking scores or using ERNIE as the feature extractor. There exists a number of alternatives to these choices. In fact, we have tried to use XGBoost [79] and Multi-Layer Perception (MLP) to replace LightGBM for ranking score regression. The performance is not as good as LightGBM in our settings. Specifically, when they achieve the best performance under the same setting, COLTR outperforms XGBoost and MLP with an average of 4.26% and 2.74% on  $NDCG_4$ . Due to the page limits, we can not include all these experiments in this section.

### 7 CONCLUSION AND DISCUSSION

In this paper, we design, implement and deploy COLTR – namely *Co-trained and Over-parameterized LTR* system at Baidu search for learning to rank tasks under semi-supervised settings. COLTR consists of three steps: (1) *RFF-based over-parameterization* enabling representation learning in the interpolating regime, (2) *Listwise-based Self-training* initializing pseudo-labels of unlabeled samples with a self-supervised listwise LTR model, and (3) *Multi-Loss Co-training for LTR* making pointwise and listwise models learn from each other. To the best of our knowledge, this work is the first to study semi-supervised training for LTR models

with labeled/unlabeled query-webpage pairs by addressing the mathematical phenomenon of interpolating in LTR tasks and the diversity of LTR models trained with various loss functions. To verify the effectiveness of COLTR, we conduct extensive offline and online experiments compared with a large number of baseline methods. Offline experiment results show that COLTR could achieve significant gain over baselines on  $\Delta NDCG_4$  under various ratios of labeled samples. Furthermore, COLTR deployed at Baidu Search significantly boosts the online ranking performance in real-world applications, which is consistent with offline results.


In this work, we actually design LTR algorithms and COLTR on top of a pre-trained language models-based retrieval method, which supplies the candidates for webpage ranking and the features of LTR. The retrieval models are end-to-end trained incorporating both self-supervision and supervision signals, such as raw text-to-text matching, click-throughs, and dwell time that characterizes the relevance between queries and webpages. We however don't include LTR models in the end-to-end training with the language models, primarily due to two reasons: (1) the requests to re-training LTR models would be more frequent to adapt to the fast shift the internet interests, while language models would be updated infrequently in a low-cost fashion; (2) given the outputs of language models (i.e., features for LTR) for either queries or webpages, their RFF transformations are restored as immediate results for online ranking. In this way, given a query and a webpage for online ranking, COLTR pickups their RFF-transformed representations and passes them to the GBM for inference in a fast manner. In the future, we attempt to study the low-cost end-to-end neural LTR approaches and their practical deployment on real-world web-scale search systems.

### REFERENCES

[1] N. Choudhary, N. Rao, K. Subbian, and C. K. Reddy, "Graph-based multilingual language model: Leveraging product relations for search relevance," in *Proceedings of the 28th ACM SIGKDD*

- Conference on Knowledge Discovery and Data Mining*, 2022, pp. 2789–2799.
- [2] X. Chu, J. Zhao, L. Zou, and D. Yin, “H-ernie: A multi-granularity pre-trained language model for web search,” in *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*, 2022, pp. 1478–1489.
- [3] Y. Sun, S. Wang, Y. Li, S. Feng, H. Tian, H. Wu, and H. Wang, “Ernie 2.0: A continual pre-training framework for language understanding,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 8968–8975.
- [4] L. Zou, S. Zhang, H. Cai, D. Ma, S. Cheng, S. Wang, D. Shi, Z. Cheng, and D. Yin, “Pre-trained language model based ranking in baidu search,” in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 4014–4022.
- [5] A. El-Kishky, T. Markovich, S. Park, C. Verma, B. Kim, R. Eskander, Y. Malkov, F. Portman, S. Samaniego, Y. Xiao, and A. Haghighi, “Twhin: Embedding the twitter heterogeneous information network for personalized recommendation,” in *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* 2022, 2022, pp. 2842–2850.
- [6] B. Pang, C. Li, Y. Liu, J. Lian, J. Zhao, H. Sun, W. Deng, X. Xie, and Q. Zhang, “Improving relevance modeling via heterogeneous behavior graph learning in bing ads,” in *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* 2022, 2022, pp. 3713–3721.
- [7] T. Yu, Y. Yang, Y. Li, X. Chen, M. Sun, and P. Li, “Combo-attention network for baidu video advertising,” in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 2474–2482.
- [8] H. Liu, Q. Gao, X. Liao, G. Chen, H. Xiong, S. Ren, G. Yang, and Z. Zha, “Lion: A gpu-accelerated online serving system for web-scale recommendation at baidu,” in *KDD '22: The 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* 2022, 2022, pp. 3388–3397.
- [9] J. Ouyang, M. Noh, Y. Wang, W. Qi, Y. Ma, C. Gu, S. Kim, K.-i. Hong, W.-K. Bae, Z. Zhao *et al.*, “Baidu kunlun an ai processor for diversified workloads,” in *2020 IEEE Hot Chips 32 Symposium*, 2020, pp. 1–18.
- [10] B. Settles, “From theories to queries: Active learning in practice,” in *Active Learning and Experimental Design workshop In conjunction with AISTATS 2010*. JMLR Workshop and Conference Proceedings, 2011, pp. 1–18.
- [11] X. J. Zhu, “Semi-supervised learning literature survey,” University of Wisconsin-Madison Department of Computer Sciences, Tech. Rep., 2005.
- [12] X. Zhu and A. B. Goldberg, “Introduction to semi-supervised learning,” *Synthesis lectures on artificial intelligence and machine learning*, vol. 3, no. 1, pp. 1–130, 2009.
- [13] Y. Chong, Y. Ding, Q. Yan, and S. Pan, “Graph-based semi-supervised learning: A review,” *Neurocomputing*, vol. 408, pp. 216–230, 2020.
- [14] A. Subramanya and P. P. Talukdar, “Graph-based semi-supervised learning,” *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 8, no. 4, pp. 1–125, 2014.
- [15] M. Bilenko, S. Basu, and R. J. Mooney, “Integrating constraints and metric learning in semi-supervised clustering,” in *Proceedings of the twenty-first international conference on Machine learning*, 2004, pp. 142–149.
- [16] Q. Xie, M.-T. Luong, E. Hovy, and Q. V. Le, “Self-training with noisy student improves imagenet classification,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 10 687–10 698.
- [17] Z. Zheng, K. Chen, G. Sun, and H. Zha, “A regression framework for learning ranking functions using relative relevance judgments,” in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007, pp. 287–294.
- [18] G. Ke, Q. Meng, T. Finley, T. Wang, W. Chen, W. Ma, Q. Ye, and T.-Y. Liu, “Lightgbm: A highly efficient gradient boosting decision tree,” *Advances in neural information processing systems*, vol. 30, pp. 3146–3154, 2017.
- [19] C. Lucchese, C. I. Muntean, F. M. Nardini, R. Perego, and S. Trani, “Rankeval: An evaluation and analysis framework for learning-to-rank solutions,” in *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2017, pp. 1281–1284.
- [20] Z. Qin, L. Yan, H. Zhuang, Y. Tay, R. K. Pasumarthi, X. Wang, M. Bendersky, and M. Najork, “Are neural rankers still outperformed by gradient boosted decision trees?” in *International Conference on Learning Representations*, 2020.
- [21] T. Qin, T.-Y. Liu, J. Xu, and H. Li, “Leter: A benchmark collection for research on learning to rank for information retrieval,” *Information Retrieval*, vol. 13, no. 4, pp. 346–374, 2010.
- [22] T. Qin and T.-Y. Liu, “Introducing letor 4.0 datasets,” *arXiv preprint arXiv:1306.2597*, 2013.
- [23] O. Chapelle and Y. Chang, “Yahoo! learning to rank challenge overview,” in *Proceedings of the learning to rank challenge*. PMLR, 2011, pp. 1–24.
- [24] T. Joachims, “Optimizing search engines using clickthrough data,” in *Proceedings of the eighth ACM SIGKDD international Conference on Knowledge Discovery & Data Mining*, 2002, pp. 133–142.
- [25] Z. Zheng, K. Chen, G. Sun, and H. Zha, “A regression framework for learning ranking functions using relative relevance judgments,” in *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, 2007, pp. 287–294.
- [26] Y. Liu, W. Lu, S. Cheng, D. Shi, S. Wang, Z. Cheng, and D. Yin, “Pre-trained language model for web-scale retrieval in baidu search,” in *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, 2021, pp. 3365–3375.
- [27] Z.-H. Zhou, M. Li *et al.*, “Semi-supervised regression with co-training,” in *IJCAI*, vol. 5, 2005, pp. 908–913.
- [28] L. Didaci and F. Roli, “Using co-training and self-training in semi-supervised multiple classifier systems,” in *Joint IAPR International Workshops on Statistical Techniques in Pattern Recognition (SPR) and Structural and Syntactic Pattern Recognition (SSPR)*. Springer, 2006, pp. 522–530.
- [29] S. Qiao, W. Shen, Z. Zhang, B. Wang, and A. Yuille, “Deep co-training for semi-supervised image recognition,” in *Proceedings of the european conference on computer vision (ECCV)*, 2018, pp. 135–152.
- [30] Z.-H. Zhou, “Ensemble learning,” in *Machine learning*. Springer, 2021, pp. 181–210.
- [31] H. Chipman, E. George, and R. McCulloch, “Bayesian ensemble learning,” *Advances in neural information processing systems*, vol. 19, 2006.
- [32] E. Raisi and B. Huang, “Co-trained ensemble models for weakly supervised cyberbullying detection,” in *NIPS Workshop on Learning with Limited Labeled Data*, 2017.
- [33] M. Belkin, D. Hsu, S. Ma, and S. Mandal, “Reconciling modern machine-learning practice and the classical bias-variance trade-off,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 32, pp. 15 849–15 854, 2019.
- [34] M. Belkin, “Fit without fear: remarkable mathematical phenomena of deep learning through the prism of interpolation,” *Acta Numerica*, vol. 30, pp. 203–248, 2021.
- [35] W. S. Cooper, F. C. Gey, and D. P. Dabney, “Probabilistic retrieval based on staged logistic regression,” in *Proceedings of the 15th annual international ACM SIGIR conference on Research and development in information retrieval*, 1992, pp. 198–210.
- [36] P. Li, Q. Wu, and C. Burges, “Mcrank: Learning to rank using multiple classification and gradient boosting,” in *Advances in Neural Information Processing Systems*, 2008, pp. 65–72.
- [37] M. Taylor, J. Guiver, S. Robertson, and T. Minka, “Softrank: optimizing non-smooth rank metrics,” in *Proceedings of the 2008 International Conference on Web Search and Data Mining*, 2008, pp. 77–86.
- [38] Z. Cao, T. Qin, T.-Y. Liu, M.-F. Tsai, and H. Li, “Learning to rank: from pairwise approach to listwise approach,” in *Proceedings of the 24th International Conference on Machine Learning*, 2007, pp. 129–136.
- [39] C. Burges, T. Shaked, E. Renshaw, A. Lazier, M. Deeds, N. Hamilton, and G. Hullender, “Learning to rank using gradient descent,” in *Proceedings of the 22nd international conference on Machine learning*, 2005, pp. 89–96.
- [40] K. Järvelin and J. Kekäläinen, “Ir evaluation methods for retrieving highly relevant documents,” in *ACM SIGIR Forum*. ACM New York, NY, USA, 2017, pp. 243–250.
- [41] S. Bruch, M. Zoghi, M. Bendersky, and M. Najork, “Revisiting approximate metric optimization in the age of deep neural networks,” in *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR 2019, 2019, pp. 1241–1244.

- [42] Z. Yuan, Z.-H. Qiu, G. Li, D. Zhu, Z. Guo, Q. Hu, B. Wang, Q. Qi, Y. Zhong, and T. Yang, "Libauc: A deep learning library for x-risk optimization," 2022.
- [43] M. Li, X. Liu, J. van de Weijer, and B. C. Raducanu, "Learning to rank for active learning: A listwise approach," in *25th International Conference on Pattern Recognition, ICPR 2020*, 2020, pp. 5587–5594.
- [44] R. Wang, R. Shivanna, D. Z. Cheng, S. Jain, D. Lin, L. Hong, and E. H. Chi, "DCN V2: improved deep & cross network and practical lessons for web-scale learning to rank systems," in *WWW '21: The Web Conference 2021*, 2021, pp. 1785–1797.
- [45] Y. Zhang, Z. Yang, B. Yu, H. Chen, Y. Li, and X. Zhao, "Structure-enhanced graph representation learning for link prediction in signed networks," in *Knowledge Science, Engineering and Management: 14th International Conference, KSEM, 2021*, pp. 40–52.
- [46] S. Guo, L. Zou, Y. Liu, W. Ye, S. Cheng, S. Wang, H. Chen, D. Yin, and Y. Chang, "Enhanced doubly robust learning for debiasing post-click conversion rate estimation," in *SIGIR '21: The 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2021, pp. 275–284.
- [47] T. Werner, "A review on instance ranking problems in statistical learning," *Mach. Learn.*, vol. 111, no. 2, pp. 415–463, 2022.
- [48] X. Wu, H. Chen, J. Zhao, L. He, D. Yin, and Y. Chang, "Unbiased learning to rank in feeds recommendation," in *The Fourteenth ACM International Conference on Web Search and Data Mining*, 2021, pp. 490–498.
- [49] A. Vardasbi, M. de Rijke, and I. Markov, "Cascade model-based propensity estimation for counterfactual learning to rank," in *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*, 2020, pp. 2089–2092.
- [50] L. Yan, Z. Qin, H. Zhuang, X. Wang, M. Bendersky, and M. Najork, "Revisiting two tower models for unbiased learning to rank," in *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2022, p. 2410–2414.
- [51] M. Szummer and E. Yilmaz, "Semi-supervised learning to rank with preference regularization," in *Proceedings of the 20th ACM international conference on Information and knowledge management*, 2011, pp. 269–278.
- [52] H. Zhuang, X. Wang, M. Bendersky, and M. Najork, "Feature transformation for neural ranking models," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020, pp. 1649–1652.
- [53] A. Rahimi and B. Recht, "Random features for large-scale kernel machines," in *Advances in Neural Information Processing Systems 20, Proceedings of the Twenty-First Annual Conference on Neural Information Processing Systems*. Curran Associates, Inc., 2007, pp. 1177–1184.
- [54] A. Blum and T. Mitchell, "Combining labeled and unlabeled data with co-training," in *Proceedings of the eleventh annual conference on Computational learning theory*, 1998, pp. 92–100.
- [55] C. Rosenberg, M. Hebert, and H. Schneiderman, "Semi-supervised self-training of object detection models," *WACV/MOTION*, 2, 2005.
- [56] M. Collins and Y. Singer, "Unsupervised models for named entity classification," in *1999 Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora*, 1999.
- [57] S. Abney, "Bootstrapping," in *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, 2002, pp. 360–367.
- [58] S. Goldman and Y. Zhou, "Enhancing supervised learning with unlabeled data," in *ICML*, 2000, pp. 327–334.
- [59] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *NAACL*, 2019, pp. 4171–4186.
- [60] Y. Sun, S. Wang, Y. Li, S. Feng, X. Chen, H. Zhang, X. Tian, D. Zhu, H. Tian, and H. Wu, "Ernie: Enhanced representation through knowledge integration," *arXiv preprint arXiv:1904.09223*, 2019.
- [61] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to information retrieval*. Cambridge University Press, 2008. [Online]. Available: <https://nlp.stanford.edu/IR-book/pdf/irbookprint.pdf>
- [62] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, 2017, pp. 5998–6008.
- [63] C. Burges, R. Ragno, and Q. Le, "Learning to rank with nonsmooth cost functions," *Advances in neural information processing systems*, vol. 19, pp. 193–200, 2006.
- [64] S. Bruch, "An alternative cross entropy loss for learning-to-rank," in *Proceedings of the Web Conference 2021*, 2021, pp. 118–126.
- [65] T. Liang and A. Rakhlin, "Just interpolate: Kernel "Ridgeless" regression can generalize," *The Annals of Statistics*, vol. 48, no. 3, pp. 1329 – 1347, 2020. [Online]. Available: <https://doi.org/10.1214/19-AOS1849>
- [66] V. Vapnik, *The nature of statistical learning theory*. Springer science & business media, 1999.
- [67] G. Yehudai and O. Shamir, "On the power and limitations of random features for understanding neural networks," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [68] A. Bodin and N. Macris, "Model, sample, and epoch-wise descents: exact solution of gradient flow in the random feature model," *Advances in Neural Information Processing Systems*, vol. 34, pp. 21 605–21 617, 2021.
- [69] P. Nakkiran, G. Kaplun, Y. Bansal, T. Yang, B. Barak, and I. Sutskever, "Deep double descent: Where bigger models and more data hurt," in *International Conference on Learning Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=B1g5sA4twr>
- [70] Z. Li, J.-F. Ton, D. Oglic, and D. Sejdinovic, "Towards a unified analysis of random fourier features," in *International conference on machine learning*. PMLR, 2019, pp. 3905–3914.
- [71] Z. Liao, R. Couillet, and M. W. Mahoney, "A random matrix analysis of random fourier features: beyond the gaussian kernel, a precise phase transition, and the corresponding double descent," *Advances in Neural Information Processing Systems*, vol. 33, pp. 13 939–13 950, 2020.
- [72] D. Holzmüller, "On the universality of the double descent peak in ridgeless regression," in *International Conference on Learning Representations*, 2021. [Online]. Available: <https://openreview.net/forum?id=0IO5VdnSAaH>
- [73] B. Sriperumbudur and Z. Szabó, "Optimal rates for random fourier features," *Advances in neural information processing systems*, vol. 28, 2015.
- [74] T. Yang, Y.-F. Li, M. Mahdavi, R. Jin, and Z.-H. Zhou, "Nyström method vs random fourier features: A theoretical and empirical comparison," *Advances in neural information processing systems*, vol. 25, 2012.
- [75] M. Samarin, V. Roth, and D. Belius, "Feature learning and random features in standard finite-width convolutional neural networks: An empirical study," in *Uncertainty in Artificial Intelligence*. PMLR, 2022, pp. 1718–1727.
- [76] Z.-H. Zhou and M. Li, "Semi-supervised learning by disagreement," *Knowledge and Information Systems*, vol. 24, pp. 415–439, 2010.
- [77] A. Qin, D. Hu, J. Liu, W. Yang, and D. Tan, "Fatman: Cost-saving and reliable archival storage based on volunteer resources," *Proceedings of the VLDB Endowment*, vol. 7, no. 13, pp. 1748–1753, 2014.
- [78] A. Qin, M. Xiao, J. Ma, D. Tan, R. Lee, and X. Zhang, "Directload: A fast web-scale index system across large regional centers," in *2019 IEEE 35th International Conference on Data Engineering (ICDE)*. IEEE, 2019, pp. 1790–1801.
- [79] T. Chen and C. Guestrin, "Xgboost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.



**Yuchen Li** is currently pursuing the Ph.D. degree with the Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China. He is also currently working as an intern at Big Data Lab, Baidu Research, Beijing, China. His research interests include machine learning, data mining, recommender system, and information retrieval.



**Haoyi Xiong** (Senior Member, IEEE) is currently a Principal Research Scientist at Big Data Lab, Baidu Research, Beijing, China, and also serves as a Graduate Faculty Scholar affiliated to University of Central Florida, Orlando FL. He received the Ph.D. degree in computer science from Télécom SudParis jointly with Université Pierre et Marie Curie, Évry, France, in 2015. From 2016 to 2018, he was a Tenure-Track Assistant Professor with the Department of Computer Science, Missouri University of Science and Technology, Rolla, MO and a Postdoc at University of Virginia, Charlottesville, VA from 2015 to 2016. His current research interests include AutoDL and ubiquitous computing. He has published more than 70 papers in top computer science conferences and journals.



**Qinzhong Wang** received the B.Eng. and M.Eng. degrees in control science and engineering from Harbin Engineering University, Harbin, China, in 2013 and 2016, and Ph.D. degree in 2021 from City University of Hong Kong, Hong Kong. Now he is a researcher at Baidu Research in Beijing, China. His research interests include computer vision, natural language processing, generative models, and searching.




**Linghe Kong** (Senior Member, IEEE) is currently a professor with the Department of Computer Science and Engineering at Shanghai Jiao Tong University, Shanghai, China. Before that, he was a postdoctoral researcher at Columbia University, McGill University, and Singapore University of Technology and Design. He received his Ph.D. degree from Shanghai Jiao Tong University 2013, Master degree from TELECOM SudParis 2007, and B. Eng. degree from Xidian University 2005. His research interests include wireless networks, big data, mobile computing, and Internet of things.




**Hao Liu** is currently an assistant professor at the Artificial Intelligence Thrust, Hong Kong University of Science and Technology (Guangzhou), Guangzhou, China. Prior to that, he was a senior research scientist at Baidu Research and a postdoctoral fellow at HKUST. He received the Ph.D. degree from the Hong Kong University of Science and Technology (HKUST), in 2017 and the B.E. degree from the South China University of Technology (SCUT), in 2012. His general research interests are in data mining, machine learning, and big data management, with a special focus on mobile analytics and urban computing. He has published prolifically in refereed journals and conference proceedings, such as TKDE, KDD, SIGIR, WWW, AAAI, and IJCAI.



**Haifang Li** received the Ph.D. degree from Academy of Mathematics and Systems Science, Chinese Academy of Sciences, Beijing, China. Now she is an algorithm engineer with Search Strategy Group, Baidu Inc., Beijing, China. Her current research interests include information retrieval, learning-to-rank, and data mining.



**Jiang Bian** (Member, IEEE) received the Ph.D. degree of Computer Engineering in University of Central Florida, Orlando, FL. He received the B.Eng. degree of Logistics Systems Engineering in Huazhong University of Science and Technology, Wuhan, China, in 2014, and the M.Sc. degree of Industrial Systems Engineering in University of Florida at Gainesville, FL, in 2016. He is currently with the Big Data Laboratory, Baidu Research, Beijing, China. His research interests



**Shuaiqiang Wang** is currently a Senior Algorithm Engineer at Baidu Inc., Beijing, China, leading the ranking strategy group that advances the document ranking for the Baidu Search Engine. Previously, he was a Research Scientist and Senior Algorithm Engineer at JD Inc., taking responsibility for the feed recommendation at JD.com. Before that, he worked as an Assistant Professor at the University of Manchester in the UK in 2017 and the University of Jyväskylä in Finland from 2014 to 2017 respectively. Earlier, he served as an Associate Professor at Shandong University of Finance and Economics in China from 2011 to 2014, and a Postdoctoral Researcher at Texas State University in the USA from 2010 to 2011. He received Ph.D. and B.Sc. in Computer Science from Shandong University in 2009 and 2004 respectively. He is broadly interested in several research areas including information retrieval, recommender systems and data mining. He published over 50 papers in leading journals and conferences.



**Guihai Chen** (Fellow, IEEE) earned his B.S. degree from Nanjing University in 1984, M.E. degree from Southeast University in 1987, and Ph.D. degree from the University of Hong Kong in 1997. He is a distinguished professor of Shanghai Jiao Tong University, Shanghai, China. He had been invited as a visiting professor by many universities including Kyushu Institute of Technology, Japan in 1998, University of Queensland, Australia in 2000, and Wayne State University, USA during September 2001 to August 2003. He has a wide range of research interests with focus on sensor network, peer-to-peer computing, high-performance computer architecture and combinatorics.



**Dejing Dou** (Senior Member, IEEE) is the Head of Big Data Lab (BDL) and Business Intelligence Lab (BIL) at Baidu Research, Beijing, China. He is also a full Professor (on leave) from the Computer and Information Science Department at the University of Oregon and has led the Advanced Integration and Mining (AIM) Lab since 2005. He has been the Director of the NSF IU-CRC Center for Big Learning (CBL) since 2018. He was a visiting associate Professor at Stanford Center for Biomedical Informatics Research during 2012-2013. Prof. Dou received his bachelor degree from Tsinghua University, China in 1996 and his Ph.D. degree from Yale University in 2004. His research areas include artificial intelligence, data mining, data integration, NLP, and health informatics. Dejing Dou has published more than 100 research papers.



**Dawei Yin** received the BS degree from Shandong University, in 2006, and the MS and PhD degrees from Lehigh University, in 2010 and 2013. He is a Senior Director of Engineering at Baidu Inc., Beijing, China. He is managing the search science team at Baidu, leading Baidu's science efforts of web search, question answering, video search, image search, news search, app search, etc.. Previously, he was senior director, managing the recommendation engineering team at JD.com between 2016 and 2020. Prior to JD.com, he was senior research manager at Yahoo Labs, leading relevance science team and in charge of Core Search Relevance of Yahoo Search. From 2007 to 2008, he was an MPhil student in The University of Hong Kong. His research interests include data mining, applied machine learning, information retrieval and recommender system. He published more than 80 research papers in premium conferences and journals, and was the recipient of WSDM2016 Best Paper Award, KDD2016 Best Paper Award, WSDM2018 Best Student Paper Award, and ICHI 2019 Best Paper Honorable Mention.