

Learning To Rank Resources with GNN

Ulugbek Ergashev
Binghamton University
Binghamton, NY, USA
uergash1@binghamton.edu

Eduard C. Dragut
Temple University
Philadelphia, PA, USA
edragut@temple.edu

Weiyi Meng
Binghamton University
Binghamton, NY, USA
meng@binghamton.edu

ABSTRACT

As the content on the Internet continues to grow, many new dynamically changing and heterogeneous sources of data constantly emerge. A conventional search engine cannot crawl and index at the same pace as the expansion of the Internet. Moreover, a large portion of the data on the Internet is not accessible to traditional search engines. Distributed Information Retrieval (DIR) is a viable solution to this as it integrates multiple shards (resources) and provides a unified access to them. Resource selection is a key component of DIR systems. There is a rich body of literature on resource selection approaches for DIR. A key limitation of the existing approaches is that they primarily use term-based statistical features and do not generally model resource-query and resource-resource relationships. In this paper, we propose a graph neural network (GNN) based approach to learning-to-rank that is capable of modeling resource-query and resource-resource relationships. Specifically, we utilize a pre-trained language model (PTLM) to obtain semantic information from queries and resources. Then, we explicitly build a heterogeneous graph to preserve structural information of query-resource relationships and employ GNN to extract structural information. In addition, the heterogeneous graph is enriched with resource-resource type of edges to further enhance the ranking accuracy. Extensive experiments on benchmark datasets show that our proposed approach is highly effective in resource selection. Our method outperforms the state-of-the-art by 6.4% to 42% on various performance metrics.

CCS CONCEPTS

• **Information systems** → **Learning to rank**; **Combination, fusion and federated search.**

KEYWORDS

selective search, resource selection, federated search

ACM Reference Format:

Ulugbek Ergashev, Eduard C. Dragut, and Weiyi Meng. 2023. Learning To Rank Resources with GNN. In *Proceedings of the ACM Web Conference 2023 (WWW '23)*, May 1–5, 2023, Austin, TX, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3543507.3583360>

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 the author(s) 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.
WWW '23, May 1–5, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-1-4503-9416-1/23/04...\$15.00
<https://doi.org/10.1145/3543507.3583360>

1 INTRODUCTION

With the World Wide Web playing an increasingly central role in our daily lives, it is constantly growing with an overwhelming amount of heterogeneous and distributed information sources. A conventional web search engine or a centralized information retrieval system cannot explore and index web documents at the same rate that they change [17, 40]. Distributed Information Retrieval (DIR) [4], also known as Federated Search (FS) [10, 33, 34, 46], has received much attention over the years as an effective solution to this problem. DIR aggregates distributed information sources (resources) under a single query interface. A basic DIR system consists of multiple resources and a broker. When a broker receives a query from a user, it passes the query to a subset of carefully selected resources to satisfy the user's information need [32, 37]. After a query is processed by the selected resources, a ranked list of documents from each resource are retrieved to the broker [52]. Finally, the broker merges the obtained results into a unified ranked list for presentation to the user [12, 13, 54]. This paper focuses on resource representation and selection in DIR.

A rich line of methods exists for the resource representation and selection problems. Recent years have seen a great interest of application of learning to rank algorithms to the resource selection problem [2, 7, 50]. However, these approaches do not consider fine-grained semantic information from both queries and resources, and are mainly limited to the statistical information like term and document frequencies of the sampled documents available in a Centralized Sample Index (CSI) [36]. Moreover, they do not consider two pieces of important information: query-resource and resource-resource relationships. Some resources may have similar or even the same documents. Those resources should have a better chance of being relevant for the same query. One such approach uses a joint probabilistic classification model for resource selection [22]. The model utilizes logistic regression to estimate the probability of relevance of information sources. While this approach is still one of the flagship solutions to the problem, it predates deep learning era and has a number of limitations, including that logistic regression requires no multicollinearity between independent variables, and its sensitivity to outliers [44].

In this paper, we propose a learning to rank based resource selection algorithm, named FedGNN. Learning to rank has attracted lots of research interest for ranking documents [15, 35, 39, 53], however, its application to ranking resources has not been studied in depth [7, 22, 50]. Resource selection is tightly coupled with the resource representation problem [46, 50]. Hence, in this paper, we propose a solution that addresses both of these problems. Specifically, we utilize a pre-trained language model (PTLM) [14] to extract semantic representations (embeddings) from query and resource data. Then, we leverage GNN to capture relationships between queries and resources. A GNN-based approach facilitates generalization across

graph which allows the model to provide an appropriate response to a new query. This inductive capability of GNN is essential for high-throughput DIR systems, which operate on evolving graphs and constantly encounter unseen queries [18].

This paper has the following contributions:

- To our knowledge, this is the first study to model a whole corpus as a heterogeneous graph for the resource selection problem. We leverage GNN that enriches embeddings with structural information from a heterogeneous graph.
- Our model extracts both (hidden) textual and structural features. Different from other studies, we represent queries and resources as embeddings from a pre-trained language model, which help captures richer semantic information.
- The empirical results and analysis demonstrate that our methods outperform state-of-the-art resource selection methods by a significant margin. Our best method outperforms the best performing baselines by 6.4% to 42% on various performance metrics.

The rest of this paper is organized as follows. Section 2 discusses related work. Section 3 presents the details of the proposed method. Section 4 describes and discusses the experimental results. Conclusion to this study and future work are given in Section 5.

2 RELATED WORK

We give an overview of the related literature in this section. We group existing resource selection methods broadly into four groups: lexicon-based, sample-based, supervised, and combination-based methods [1, 26, 46].

Lexicon-based methods treat each resource as a big bag of words, and they collect word statistics from all documents to describe a resource [4, 51]. These approaches calculate the similarity between user query and each resource based on matching term frequencies, adapting traditional document ranking methods [3]. This group of approaches has the limitation of not being able to encode semantics both at word and sentence granularity, which reduces the resource representation quality [19].

Sample-based algorithms sample documents from each resource and build a Centralized Sample Index (CSI). Then, a user query is issued to the CSI to retrieve relevant documents, and each relevant document sampled from a resource is considered as a vote for the resource. Many methods of in this group, such as ReDDE [47], CRCS [45] and SUSHI [48], utilize various voting algorithms over documents to infer a ranking of the resources. KBCS [19], the most recent method in this group, represents resources as a weighted entity set based on context- and structure-based measures, and ranks resources according to a query-resource similarity metric. Both lexicon-based and sample-based methods have been proved effective when the resources are mostly homogeneous [26], but they cannot fully utilize structural features of text inputs.

Supervised methods employ machine learning approaches to build a model to tackle the resource selection problem. There are three types of supervised methods: query classification, resource classification and learning to rank methods. Dai et al. [7] proposed a learning to rank method where they trained an SVMrank model based on features such as query-independent information, term-based statistics and CSI-based features. Another method is LTRRS [50], which uses term matching, CSI based, and topical relevance

features. These features are combined to build a LambdaMART model that directly optimizes resources ranking list based on the nDCG metric to rank resources. The main drawbacks of these group of methods are that they require extensive feature engineering and heavily rely on domain knowledge, which makes it difficult to capture diverse feature aspects of queries and resources.

Combination-based methods combine some of the approaches mentioned above. The method ECOMSVZ [24] combines a number of strategies to rank resources according to a given query, namely i) topical closeness between a query and resources, ii) estimation of a resource popularity from online sources, and iii) query expansion. The authors of [49] present a method that aggregates all three strategies proposed in [24] and ranks resources by giving more weights to small (specialized) resources.

The work most related to ours is that of Hong et al [22]. It proposes a joint probabilistic classification model for resource selection. The model utilizes logistic regression to estimate the probability of relevance of information sources in a joint manner by considering both the evidence of individual sources and their relationships. However, logistic regression has a number of assumptions and limitations for modeling relations between variables [41]. Some limitations include considering dependent variable structure, observation independence, absence of multicollinearity, linearity of independent variables and log odds, and large sample size [44]. In practice, if the data used do not satisfy the assumed conditions, the use of logistic regression model may incur significant errors [20].

We propose a new learning to rank approach that utilizes a PTLM and GNN to address the limitations listed above. Our method takes into account both the semantics of query and resource as well as the query-resource and resource-resource relationships. First, we obtain embeddings from a pre-trained language model for queries and resources. Then, we construct a heterogeneous graph based on query-resource relevance judgement scores. To our knowledge, our approach is the first to employ a pre-trained language model BERT and GNN for the resource selection problem. This gives us a richer resource representation, which in turn allows us to identify the most relevant resources for a given user query.

3 OUR METHOD

In this section, we describe FedGNN. We start with the task formulation. The objective of the resource selection task is to output a ranked list of resources given a query representing an information need. Our approach falls into the learning to rank class of methods.

Let $R = \{R_1, R_2, \dots, R_N\}$ be the resources to be ranked for a given user query q . Suppose the relevance scores of the resources are given as multi-level ratings $Z = \{l(R_1, q), l(R_2, q), \dots, l(R_N, q)\}$, where $l(R_i, q) \in \{r_1, \dots, r_N\}$ denotes the relevance score of R_i for a user query q . Without loss of generality, we assume $l(R_i, q) \in \{0, 1, \dots, N - 1\}$ and name the corresponding relevance score as N -level ratings. If $l(R_i, q) > l(R_j, q)$, then R_i should be ranked before R_j . Let \mathcal{F} be the function class and $g \in \mathcal{F}$ be a ranking function. The goal of learning to rank is to learn the optimal ranking function from the training data by minimizing a certain loss function defined on the resources and their labels.

Since we do not know the optimal g , we use θ to denote the parameters that we need to learn toward finding (an approximation

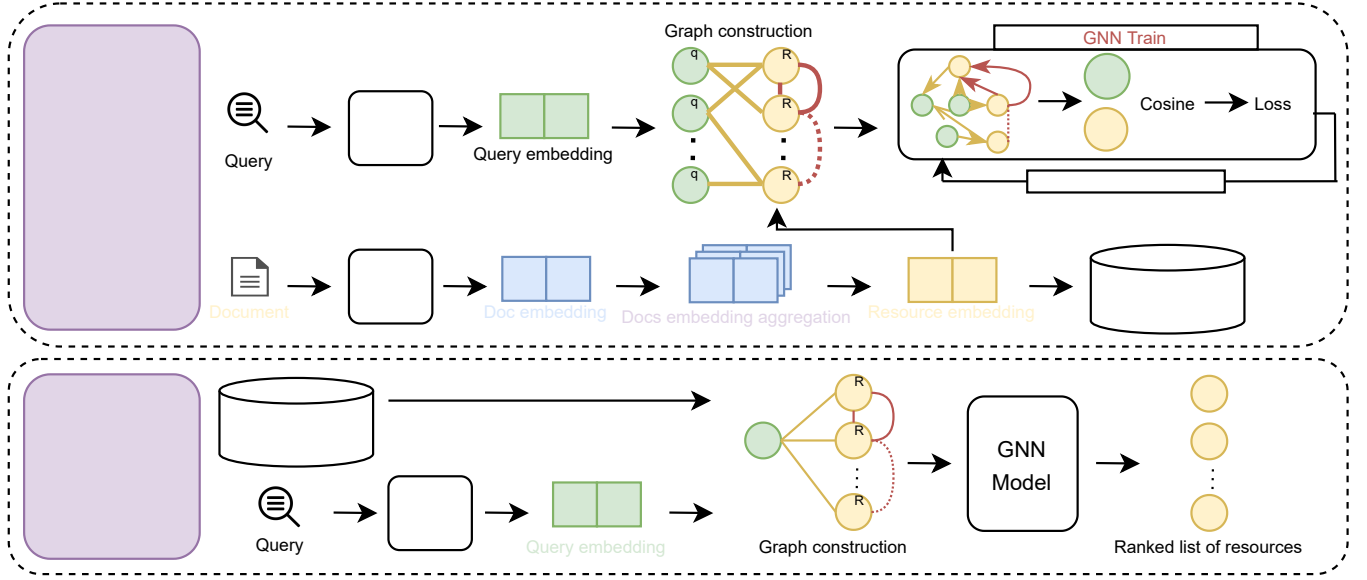


Figure 1: The overall architecture of the proposed method. Documents and queries are fed into a BERT encoder one at a time for semantic feature extraction. After obtaining resource embeddings from the aggregated document embeddings, resource embeddings are stored for the usage in the inference phase. A heterogeneous graph is constructed using query-document pairs where the edges in yellow and red denote query-resource and resource-resource edges, respectively. GNN is trained based on the constructed graph. In the inference phase, after obtaining an embedding from a BERT encoder for a user query, a graph is built based on the stored resource embeddings. Finally, a trained GNN model outputs a ranked listed of resources.

of) g . We propose to learn the optimal parameters θ of a function $g(\cdot; \theta)$ from the training data by minimizing a pointwise loss function defined as

$$\mathcal{L}(R, q; \theta) = \sum_{i=1}^n (\text{Cosine}(g(R_i, q; \theta)) - l(R_i, q))^2 \quad (1)$$

where g represents the GNN model.

There is a counter-intuitive connection between ranking measures and loss functions in learning to rank methods. On the one hand, we learn a ranking function by minimizing a pointwise loss function. On the other hand, it is the ranking measures, like nDCG or nP@k, that we use to evaluate the performance of the learned ranking function. Thus, one may wonder if the minimization of the loss functions can lead to the optimization of the ranking measures. A number of works looked into this question and proved that the minimization of these loss functions will lead to the maximization of the ranking measures [5, 6, 31]. Relevant to our work, it has been proved that the regression and classification based losses used in the pointwise approach are upper bounds of $(1 - \text{nDCG})$ [6, 31].

3.1 Architecture Overview

Figure 1 depicts the architecture of our proposed method. The method consists of the following five modules: (i) Dataset preprocessing, (ii) BERT encoder based feature extraction, (iii) Resource embedding generation from the aggregated documents embeddings, (iv) Graph construction, and (v) GNN based feature extraction with similarity score prediction.

Given a collection of search results (query-document pairs) with human labelled relevance judgement, we select top N documents per resource with the highest labelled judgements. Then, we input documents and queries one at a time into a BERT encoder for semantic feature extraction. After obtaining resource embeddings from the aggregated document embeddings, we store resource embeddings for the usage in the inference phase, and also we construct a heterogeneous graph according to the provided query-document pairs. In the training phase, the constructed heterogeneous graph is fed into GNN to extract the structural features for query and resource nodes. Then, we minimize Mean Squared Error (MSE) based on relevance judgement and computed Cosine similarity score between query and resource. In the inference phase, we obtain an embedding from a BERT encoder for a user query and build a graph based on the stored resource embeddings so that a query node connects to each resource node. Finally, a trained GNN model outputs a similarity score for each resource.

3.2 Data Preprocessing

The Data Preprocessing module is mainly responsible for the selection of sampled documents to represent each resource. Instead of considering all the sampled documents for every training query, the module takes only a subset of top-ranked documents over all queries. Specifically, we sort the documents based on their similarities with training queries, and select the N unique documents with the overall highest similarities, per each resource. Then, we use a BERT encoder to obtain semantic representations for queries and top N documents for each resource.

3.3 Query and Resource Representations

To represent a document, we use its title concatenated with its body text, while to represent a query, we use query text. We adopt a feature-based approach to obtain query and resource representations using a BERT encoder. There are major computational benefits to pre-compute an expensive representation of the training data once and then run many experiments with cheaper task-specific learning to rank models on top of this representation [11]. Hence, we first construct a heterogeneous graph from the representations and store resource embeddings for the further usage in the inference phase, then we train a model with a GNN algorithm.

BERT encoder takes a sequence of tokens as input and outputs a sequence of their corresponding final embeddings, i.e., one embedding for each token. Different from other feature-based approaches [30], we do not add special token [CLS] at the beginning of the token sequence. Following the original BERT [11, 16], each query and document is converted to a token sequence using the BERT tokenizer, as in Equation 2.

$$\{e_1, \dots, e_i\} = \text{BertTokenizer}(\{t_1, \dots, t_i\}) \quad (2)$$

where $\{e_1, \dots, e_i\}$ and $\{t_1, \dots, t_i\}$ represent a token sequence and the words in a text (i.e., query or document), respectively. Then, we feed the token sequence into a BERT encoder to obtain a token representation

$$\{E_1, \dots, E_i\} = \text{BERT}(\{e_1, \dots, e_i\}) \quad (3)$$

where $\{E_1, \dots, E_i\}$ is the last hidden state of BERT output. The last hidden state contains one embedding for each token (in other words, for each word in a text) where the size of a single embedding is 768. To obtain a final representation for an input (query or document), we calculate the mean of token embeddings.

$$H = \frac{1}{K} \sum_{i=1}^K E_i \quad (4)$$

where K is the number of tokens in a sequence. Formally, we denote feature representations for queries and documents as $H_i^Q \in R^d$ and $H_i^D \in R^d$, where d is the embedding dimensionality, while q and D denote query and document embeddings, respectively. Once the feature embeddings for the selected top N documents are obtained, feature embedding for the j -th resource is calculated by taking the arithmetic mean of the top N documents per each resource:

$$H_j^R = \frac{1}{N} \sum_{i=1}^N H_i^D \quad (5)$$

3.4 Graph Construction

We build a heterogeneous graph, denoted as $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{R})$, where \mathcal{V} and \mathcal{E} are a set of nodes and edges, respectively. \mathcal{R} represents a set of edge relations. For example, a tuple $\langle q_i, \text{query-resource}, R_j \rangle \in \mathcal{E}$ denotes that a query node q_i and a resource node R_j are connected by the relation $\text{query-resource} \in \mathcal{R}$. There are two types of edge relations $\mathcal{R} = \{\text{query-resource}, \text{resource-resource}\}$. For convenience, we denote query-resource and resource-resource edge relations as qr and rr , respectively. Query and resource node embeddings are initialized with the output from the BERT encoder.

From the given *query-document* relevance judgement scores, we calculate the *query-resource* relevance score by summing the *query-document* relevance scores (grouped by query and resource) multiplied by a normalization coefficient α :

$$\hat{z}_{ij} = \alpha * \sum_{k=1}^n r_k \quad (6)$$

where \hat{z}_{ij} is the ground truth relevance score between the i -th query and the j -th resource, r_k is the ground truth relevance judgement score between the query and the k -th document, and n denotes the number of data points. We also utilize ground truth relevance scores as the weights on edges of the *query-resource* relation type. For convenience, we use E^{qr} to denote weights on edges of the *query-resource* relation type.

Weights on edges of the *resource-resource* relation type are derived from the Cosine similarity function between two resource embeddings:

$$E_{ij}^{rr} = \text{Cosine}(H_i^R, H_j^R) \quad (7)$$

where H_i^R and H_j^R are embeddings of the i -th and the j -th resources, respectively. Figure 1 illustrates a constructed heterogeneous graph, where yellow lines represent *query-resource* edges and red lines are of the *resource-resource* edge type.

3.5 The GNN Module

In this paper, we adopt Relational Graph Convolutional Network (R-GCN) [43] as a representative from the GNN family because it is designed to cope with heterogeneous type of graphs. We use R-GCN to capture the structural nature of the resource selection problem. Next, we will briefly introduce R-GCN.

3.5.1 Relational Graph Convolutional Network. R-GCN [43] is an extension of the Graph Convolutional Network (GCN), which performs the information exchange in a multi-relational graph. As discussed in Section 3.3, we leverage BERT encoder to initialize H^0 for query and resource nodes. The representation of a node for the l -th layer is defined by:

$$\bar{H}_i^{(l+1)} = \sigma \left(\sum_{r \in \mathcal{R}} \sum_{j \in N_i^r} E_{ij}^r W_r^l H_j^l + W_0^l H_i^l \right) \quad (8)$$

where $\mathcal{R} = \{qr, rr\}$ is the set of edge relation types and N_i^r denotes the set of neighbors of node v_i based on edges of relation type $r \in \mathcal{R}$. W_r^l and W_0^l are learnable parameter matrices, and E_{ij}^r is a scalar weight on the edge of type r between the i -th and the j -th nodes. To simplify, we can write such a message passing process as:

$$\bar{H}^{(l+1)} = \text{GNN}(H^l; \theta) \quad (9)$$

where θ is the learning parameter and \bar{H} is the final representations for query and resource nodes.

3.6 Final Prediction and Loss Function

To predict a relevance score between a query and a resource, we employ the Cosine similarity:

$$z_{ij} = \text{Cosine}(\bar{H}_i^Q, \bar{H}_j^R) \quad (10)$$

where \bar{H}_i^q and \bar{H}_j^R are the final embedding vectors for the i -th query q_i and the j -th resource R_j obtained from GNN, respectively.

Finally, we adopt the pointwise ranking objective as a loss function. The idea of pointwise ranking function is to measure the distance between the ground truth relevance score and the predicted relevance score. For this, we utilize Mean Square Error (MSE) loss:

$$\mathcal{L} = \sum_{i=1}^n (\hat{z}_{ij} - z_{ij})^2 \quad (11)$$

where z_{ij} and \hat{z}_{ij} are predicted and ground truth relevance scores, respectively, and n denotes the number of data points.

4 EXPERIMENTS

To evaluate the effectiveness of our proposed method in resource selection, in this section, we validate this method and compare it against a number of baseline models. First, we describe baseline models and real-world public datasets used in our comprehensive experimental study and provide the experimental details. Finally, we present empirical results with discussions and analysis.

4.1 Baselines

To examine the performance of the proposed model, we take all four categories of resource selection approaches as baselines, including lexicon-based (Taily [3]), sample-based (Rank-S [28], ReDDE [47], CRCS [45] and KBCS [19]), supervised (L2R [7] and Jnt [22]) and combination-based (SSLTS [49] and ECOMSVZ [24]). These baseline methods are briefly described below:

KBCS [19] is one of the recently proposed methods in the sample-based category that has the best performance among methods of this category.

L2R [7] is a learning to rank method that trains SVMrank based on features such as query-independent information, term-based statistics and sample-document features.

ECOMSVZ [24] is combination-based method that achieves one of the best results for selecting relevant resources in TREC 2014. It combines a number of strategies to rank resources for a given query, namely i) topical closeness between a query and resources, ii) estimation of a resource popularity from online sources, and iii) query expansion.

SSLTS [49] is also a combination-based method and consists of all three strategies proposed in [24] and ranks resources by giving more weights to small (specialized) resources.

Taily [3] is a method in the lexicon-based category. It models a query's score distribution in each resource as a Gamma distribution and selects resources with highly scored documents in the tail of the distribution.

Rank-S [28] is a sample-based algorithm from the SHiRE family of methods, where the scores of retrieved documents from the CSI are decayed exponentially and then treated as votes for the resources the documents were sampled from.

ReDDE [47] is another effective and commonly used sample-based resource selection method.

CRCS [45] is a sample-based method and, like ReDDE, it runs the query on a CSI and, in contrast to ReDDE, ranks resources according to the positions of retrieved documents in the CSI ranking.

Jnt [22] is a learning to rank model. It employs a joint probabilistic

classification model that estimates the probabilities of relevance in a joint manner by considering relationships among resources.

4.2 Comparison Methodology

To compare with the baselines fairly, we compare the experimental results of our method with those that were reported in the literature directly based on the same benchmark datasets and the same evaluation metrics. More specifically, if the reported experimental result for a baseline B is based on dataset D and evaluation metric M, we compare with B using the same D and the same M. An advantage of this comparison approach is that we do not need to implement the baselines ourselves. The down side of this approach is that since different datasets and different evaluation metrics were used in reported literature, the comparison is less than uniform, i.e., we are not able to compare with all baselines using the same datasets and the same evaluation metrics uniformly.

4.3 Evaluation Metrics

As mentioned earlier, to facilitate fair comparison, we adopt the same evaluation metrics used in the baseline methods. The following evaluation metrics will be used in our experimental study:

- Precision at k or $P@k$. Precision is the number of relevant retrieved items with respect to the total number of retrieved items. $P@10$ is used to compare with baselines L2R [7], KBCS [19], ReDDE [47], CRCS [45], Rank-S [28], Taily [3] and Jnt [22].
- Normalized precision at k or $nP@k$. This metric measures the relevance scores of the top ranked k resources, normalized by the relevance scores of the best possible k resources for the given query. $nP@{1,5}$ is used to compare with baselines SSLTS [49] and ECOMSVZ [24].
- Normalized Discounted Cumulated Gain at k or $nDCG@k$. This metric allows different degrees of relevance of resources to be considered. $nDCG$ also discounts the value of a relevant resource when a resource is ranked lower. $nDCG@{10,20}$ is used to compare with SSLTS [49] and ECOMSVZ [24].

$P@k$ and $nDCG@k$ are widely used evaluation metrics in information retrieval and resource selection. The precision-oriented metrics at lower ranks such 10 or 30 (e.g., $P@10$ and $nDCG@30$) are expected to perform better than recall-based metrics at deeper ranks (e.g., $P@100$ and $nDCG@100$) in the resource selection problem because resource selection returns a small portion of resources [26]. Moreover, since users are usually interested in the first 10 results, shallow ranked relevant results are expected to directly meet users' satisfaction.

4.4 Datasets

We run our experiments on the following two widely used benchmark datasets. The baselines were evaluated on those datasets, too. **ClueWeb09-B**¹ is a subset of the full web collection ClueWeb09. The dataset was created to support research on information retrieval and related human language technologies, and it contains around 50 million English webpages. This collection is divided into different resources to evaluate resource selection methods (see below for more details).

¹<http://www.lemurproject.org/clueweb09>

Table 1: Dataset statistics

Dataset	# of resources	Total # of docs (K)	Resource (K)			Compared with baselines
			Min	Max	Avg	
ClueWeb09-B	100	50,220	63	1,691	502	KBCS, ReDDE, CRCS, Rank-S
ClueWeb09-B	123	50,220	32	734	408	L2R, ReDDE, Rank-S, Taily, Jnt
FedWeb14	149	187.7	.022	2.7	1.2	SSLTS, ECOMSVZ

Table 2: Query set statistics

Query set	# of queries	Folds	Compared with baselines
Web Track 2009	50	5	KBCS, ReDDE, CRCS, Rank-S
Web Track 2009-2012	200	10	L2R, ReDDE, Rank-S, Taily, Jnt
FedWeb14	50	5	SSLTS, ECOMSVZ

FedWeb Greatest Hits [9] is a large test collection designed to support research in web information retrieval including resource selection in federated search. The collection contains two datasets FedWeb13 and FedWeb14, which consist of 157 and 149 resources, respectively. We use the FedWeb14 dataset because it has about 3 million sampled documents. The dataset also has 275 information needs (queries), and out of those only 50 queries are judged.

According to L2R [7], experiments were conducted with the ClueWeb09-B dataset which was divided into 123 resources. The relevance scores were obtained from TREC Web Track 2009-2012². Another baseline that used the same dataset in experiments is KBCS [19]. They divided it into 100 resources, and evaluation queries were acquired from TREC Web Track 2009. Other two baseline methods SSLTS [49] and ECOMSVZ [24] used the FedWeb14 dataset, and for evaluation, judged queries were utilized.

Since our proposed learning to rank algorithm requires ground truth relevance scores to train a model, we leverage only judged portion of the queries from the FedWeb14 dataset for both training and evaluation. For the ClueWeb09-B dataset with 123 resources setup, we use TREC Web Track 2009-2012 for training and evaluation. TREC Web Track 2009 is used for the ClueWeb09-B dataset with 100 resources setup to train and evaluate a model.

Table 1 summarizes the statistics of the datasets. For the FedWeb14 dataset, we consider only judged documents. In addition, the last column “Comparison with baselines” lists the baseline methods that are used to compare with our method using a given dataset.

Table 2 gives query sets that are used to train and evaluate the proposed model. It also shows cross-validation settings we followed for each query set. In cross-validation, we partition queries. For example, FedWeb14 has 50 queries. For five-fold cross-validation, the 50 queries are partitioned into 5 subsets; in each run, four subsets are used for training and one subset is used for testing. The results are then averaged over the five runs. Graph statistics for three datasets are presented in Table 3

²<https://trec.nist.gov/data/webmain.html>**Table 3: Graphs statistics for three datasets**

Datasets	# of resource nodes	# of query nodes	# of query-resource edges	# of resource-resource edges
ClueWeb09-B	123	200	4046	7503
ClueWeb09-B	100	50	1618	4950
FedWeb14	149	50	4286	11026

Table 4: Search accuracy comparison for ClueWeb09-B with 123 resources

Methods	P@10	
	T=4	T=8
ReDDE	0.355	0.363
Rank-S	0.350	0.259
Taily	0.346	0.260
Jnt	0.370	0.367
L2R	0.374	0.377
FedGNN	0.398	0.407

Table 5: Search accuracy comparison for ClueWeb09-B with 100 resources

Methods	P@10	
	T=4	T=6
ReDDE	0.244	0.252
CRCS	0.236	0.252
Rank-S	0.316	0.365
KBCS	0.370	0.400
FedGNN	0.418	0.429

4.5 Implementation Details

To obtain the initial embeddings for document titles and queries, we utilize $BERT_{base}$ architecture that encodes input text to a 768-dimensional dense vector space. After aggregating the top 10 documents’ embeddings to represent resource nodes, we leverage the DGL³ python package to build a graph and train a GNN model. We implemented the proposed method in PyTorch. We trained FedGNN model with Adam optimizer [27] and followed cross-validation settings shown in Table 2. The results reported in this paper are based on the following parameter settings: RelGraphConv aggregator: pool, HeteroGraphConv aggregator: mean, hidden dropout rate: 0.3 (FedWeb14), input feature size: 768, hidden size: 768, output size: 768. All experiments were conducted on a Debian server equipped with CPU: Intel(R) Core(TM) i9-10940X, GPU: NVIDIA Quadro RTX 6000 24GB, RAM: 64GB. To simulate a distributed information retrieval environment, we followed same search engine setup parameters given in the baseline papers. When resource selection decides which resource to retrieve, then for each query, top 1000 documents are returned from each resource and merged directly.

4.6 Experimental Results


There are three sets of experiments to verify the effectiveness of our method. One set compares ours with L2R [7] and its baselines Taily [3], Rank-S [28], ReDDE [47], and Jnt [22] on the document level as shown in Table 4. Table 5 reports the results of the second set of experiments, showing a comparison with KBCS [19] as well as its baselines ReDDE [47], CRCS [45], and Rank-S [28] on the document level. The results of the final set of experiments are reported in Table 6, comparing the effectiveness of our method with those of SSLTS [49] and ECOMSVZ [24].

The empirical results show that FedGNN outperforms the baselines significantly in all metrics suggesting that the combination of a BERT and GNN based method can improve the effectiveness

³<https://www.dgl.ai/>

Table 6: Search accuracy comparison for FedWeb14

Methods	nDCG@10	nDCG@20	nP@1	nP@5
SSLTS (Market share)	0.55	0.66	0.38	0.48
ECOMSVZ	0.712	0.624	0.535	0.604
FedGNN	0.794	0.821	0.76	0.788

**Figure 2: The nDCG@10 measure of FedGNN for different λ**

of resource selection. Specifically, according to the results given in Table 4, our model outperforms the best baseline in its group, L2R, by 6.4% and 8% on P@10 for top 4 and 8 resources, respectively. On dataset ClueWeb09-B with 100 resources (reported in Table 5), FedGNN outperforms the best baseline in its group, KBCS, by 13% and 7.3% on nP@10 for top 4 and 6 resources, respectively. The final experiment with the dataset FedWeb14, the results provided in Table 6 show that FedGNN outperforms the best baseline in its group, ECOMSVZ, by 11.5% on nDCG@10, by 42% on nP@1, and by 30.4% on nP@5. On nDCG@20, the best baseline is SSLTS, and FedGNN outperforms this baseline by 24.4%.

One of the reasons that may have contributed to the above significant improvements is that our proposed method does not suffer from the "bigger gets bigger vote" problem. Specifically, Shokouhi et al [45] reports that the bigger the collection size, the higher the chances are that it contains more relevant documents. Since sample-based methods are based on a centralized sample index (CSI), most of those methods are biased towards bigger resources by giving larger resources higher scores. On the other hand, our proposed method does not require CSI which is expensive to obtain. Our method needs a small number of documents from each resource, and it aggregates embeddings of documents to obtain a resource embedding. Thus, we leverage a single source of truth to represent a resource. In this way, resource discrimination by its size is avoided which may have contributed to the better effectiveness of our proposed method over the baselines.

Another reason for the superiority of FedGNN is the interactive aggregation of features from BERT and GNN modules. It leads to the combination of the advantages of these two features to the greatest extent which guarantees a better effectiveness over the existing methods. More specifically, the baselines L2R, KBCS, SSLTS and ECOMSVZ utilize many term-based statistical features which do not capture semantic information. It is well-known that a user

Table 7: Ablation study of the effectiveness of the GNN

Methods	FedWeb14		ClueWeb09B-100		ClueWeb09B-123	
	nDCG@10	nP@10	nDCG@10	nP@10	nDCG@10	nP@10
FedBERT	0.088	0.107	0.261	0.405	0.343	0.575
FedHGT	0.71	0.668	0.592	0.812	0.38	0.637
FedGNN	0.794	0.808	0.631	0.856	0.561	0.788

query may have many meanings in different contexts, and hence, capturing semantic information is crucial in federated search. Different from the baselines, our model can not only consider semantic information of a user query, but also integrates the structural information between query and resource nodes. Query and resource nodes in a constructed graph get useful messages from their two-hop neighbors, and it leads to better ranking results [8, 21, 42].

4.7 Ablation Study

To analyze the usefulness of each component of FedGNN, we perform the following experiments.

4.7.1 Effectiveness of the GNN. Our model leverage the aggregated features from BERT and GNN. Compared with the original BERT, FedGNN integrates structural features in addition to semantic features. To validate the effectiveness of GNN, we design an experiment to compare FedGNN with a model without GNN. Essentially, the model employs Cosine similarity function to get a relevance score between query and resource embeddings obtained from BERT encoder. We name the model as FedBERT. Moreover, to validate the effectiveness of Relational Graph Convolutional Network (R-GCN) [43] used in our model, we substitute that with the Heterogeneous Graph Transformer (HGT) [23] algorithm; we name it FedHGT. Table 7 illustrates the experimental results. Compared with using BERT alone, FedGNN achieves significantly better results on the three datasets. This confirms the necessity of adding GNN to capture structural features. The empirical results also elucidate that R-GCN performs better than HGT, which implies R-GCN generalizes better on unseen queries in the resource selection problem. This is the chief contribution of our paper.

4.7.2 The impact of resource-resource edges. We investigate the effect of resource-resource edges on the ranking. To obtain resource-resource edges, we calculate Cosine similarity between resource nodes, and conduct an experiment to determine the optimal value of parameter λ . λ is a threshold that controls the number of resource-resource edges based on their similarity scores. For example, $\lambda = 0.5$ means we keep the edges with similarity scores at least 0.5.

In the experiment, we increase the value of λ from 0 to 1 with a fixed step of 0.1. Figure 2 illustrates the empirical results from the experiment conducted on the FedGNN model with the FedWeb14, ClueWeb09B-100 and ClueWeb09B-123 datasets. Note that when $\lambda = 1$, it will essentially exclude all resource-resource edges, leaving the graph consisting of only query-resource type of edges.

The experimental results indicate that the presence of resource-resource type of edges in a graph increases the overall ranking performance, which proves that structure-based semantic similarities are important in the resource selection problem. On FedWeb14 dataset, nDCG@10 is increased by 4% when $\lambda = 0.0$ is compared to

Table 8: Ablation study of the effect of BERT encoder

Methods	FedWeb14		ClueWeb09B-100		ClueWeb09B-123		Average	
	nDCG@10	nP@10	nDCG@10	nP@10	nDCG@10	nP@10	nDCG@10	nP@10
FedGNN-WV	0.789	0.812	0.57	0.721	0.551	0.782	0.637	0.772
FedGNN-DV	0.738	0.781	0.593	0.807	0.542	0.773	0.624	0.787
FedGNN-FT	0.806	0.812	0.577	0.75	0.547	0.751	0.643	0.771
FedGNN	0.794	0.808	0.631	0.856	0.561	0.788	0.662	0.817

$\lambda = 1.0$. While on datasets ClueWeb09B-100 and ClueWeb09B-123, nDCG@10 is increased by 3.8% and 2.4%, respectively. Moreover, the evaluation metric nDCG is maximized at the same value of parameter λ for all three datasets: $\lambda = 0.0$. Resource-resource relationships, at such a value of λ when the evaluation metric is maximized, are important and may carry crucial insights into an overall structure of a graph. Hence, λ is set to the above identified value (i.e., 0.0) in our all analysis with the baselines.

4.7.3 Effect of BERT encoder. In this section, we explore the effect of BERT encoder to the overall ranking. To validate, we substitute BERT encoder with some of the most popular language models like Word2Vec [38], Doc2Vec [29], FastText [25], and compare them with our model. Specifically, we pre-train Word2Vec, Doc2Vec and FastText language models for each dataset (e.g. FedWeb14, ClueWeb09B-100, ClueWeb09B-123). Then, we develop three models namely FedGNN-WV, FedGNN-DV, FedGNN-FT that utilize Word2Vec, Doc2Vec and FastText, respectively. Since Word2Vec and FastText generate an embedding on a word level, we average word embeddings to obtain a representation on a document level.

Table 8 shows that Word2Vec and FastText perform slightly better than BERT on nP@10 metric over FedWeb14 dataset. FastText also outperforms BERT based model only by 1.5% on nDCG@10. However, BERT outperforms the best baselines on average over three datasets by 3% and 4% on nDCG@10 and nP@10, respectively. The reasons for the superiority of BERT based model is its ability in representing the contextualized embedding. Different from other language models, BERT considers the order of words in a sentence which leads to syntactic and semantic understanding of a sentence [11]. This validates the importance of BERT encoder in our model.

4.7.4 Effect of different hidden layers of BERT encoder. $BERT_{BASE}$ architecture has 12 transformer blocks (a.k.a hidden layers or layers). Authors of BERT [11] consider some methods that generate embeddings from different hidden layers. In this experiment, we aim to determine a method that increases an overall ranking of resources by generating a better contextualized representation for queries and resources. The following methods are considered in the experiment: [CLS] token of the last hidden layer (CLS), concatenation of the last 4 hidden layers (C4LH), sum of the last 4 hidden layers (S4LH), last hidden layer (LHCLS), last hidden layer without special tokens like [CLS] (LH-CLS).

Table 9 shows the FedGNN results trained on five embedding generation methods. The results are reported on nDCG@10 metric on three datasets. The results indicate that the LH-CLS method outperforms the best methods by 1.29% on average. Hence, we utilize an embedding obtained from the last hidden layer of BERT encoder without special tokens.

Table 9: Ablation study of the effect of different hidden layers of BERT encoder on nDCG@10 metric

Datasets	CLS	C4LH	S4LH	LHCLS	LH-CLS
FedWeb14	0.773	0.742	0.775	0.778	0.794
ClueWeb09B-100	0.62	0.617	0.616	0.614	0.631
ClueWeb09B-123	0.547	0.559	0.537	0.56	0.561

Table 10: Ablation study on document title and body text

Datasets	Title		Body		Title+Body	
	nDCG@10	nP@10	nDCG@10	nP@10	nDCG@10	nP@10
FedWeb14	0.779	0.796	0.775	0.783	0.794	0.808
ClueWeb09B-100	0.617	0.833	0.62	0.82	0.631	0.856
ClueWeb09B-123	0.558	0.776	0.548	0.774	0.561	0.788

4.7.5 Effect of document title and body text. We designed an experiment to validate that the concatenation of title and body text gives a better representation of a document that leads to a higher overall ranking of resources.

Table 10 shows empirical results of three FedGNN models trained on three different document attributes: Title, Body and Title+Body. Input texts larger than 512 tokens are truncated. The experimental results show that a model trained on Title+Body attribute outperforms the best results trained on either Title or Body on average by 1.4% and 1.94% on nDCG@10 and nP@10, respectively.

5 CONCLUSION

In this paper, we tackled the resource selection problem in federated search. We proposed a learning to rank with GNN approach that addresses resource representation and resource selection problems at the same time. The proposed approach takes advantage of both a pre-trained language model and a heterogeneous graph for learning and ranking resources based on the given query. Our approach consists of five modules: data processing module, a pre-trained language module for extracting semantic information from queries and documents, a module for obtaining a resource representation, a heterogeneous graph construction module and a GNN module for capturing structural dependencies between queries and resources as well as among resources. Comprehensive experiments on two benchmark datasets demonstrate the superiority of the proposed method compared with the state-of-the-art approaches. Specifically, our method outperforms the state-of-the-art by 6.4% to 42% on various performance metrics. Moreover, the ablation studies prove the effectiveness of each module of the FedGNN model.

ACKNOWLEDGMENTS

This work was supported in part by the U.S. NSF 1838145 and 2107213 grants, Temple University office of the Vice President for Research 2022 Catalytic Collaborative Research Initiative Program and a gift from NVIDIA Corporation.

REFERENCES

- [1] Qingyao Ai, Jiaxin Mao, Yiqun Liu, and W. Bruce Croft. 2018. Unbiased Learning to Rank: Theory and Practice. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management (Torino, Italy) (CIKM '18)*. Association for Computing Machinery, New York, NY, USA, 2305–2306. <https://doi.org/10.1145/3269206.3274274>
- [2] Rami S. Alkhalaf, Deepak P., Joemon M. Jose, and Fajie Yuan. 2017. LTRo: Learning to Route Queries in Clustered P2P IR. In *Advances in Information Retrieval*, Joemon M Jose, Claudia Hauff, Ismail Sengor Altingovde, Dawei Song, Dyaa Albakour, Stuart Watt, and John Tait (Eds.). Springer International Publishing, Cham, 513–519.
- [3] Robin Aly, Djoerd Hiemstra, and Thomas Demeester. 2013. Taily: shard selection using the tail of score distributions. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2013*. Association for Computing Machinery (ACM), United States, 673–682. <https://doi.org/10.1145/2484028.2484033>; null; Conference date: 29-07-2013 Through 01-08-2013.
- [4] Jamie Callan. 2000. *Distributed Information Retrieval*. Springer US, Boston, MA, 127–150. https://doi.org/10.1007/0-306-47019-5_5
- [5] Wei Chen, Tie-yan Liu, Yanyan Lan, Zhi-ming Ma, and Hang Li. 2009. Ranking Measures and Loss Functions in Learning to Rank. In *Advances in Neural Information Processing Systems*, Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, and A. Culotta (Eds.), Vol. 22. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2009/file/2f55707d4193dc27118a0f19a1985716-Paper.pdf>
- [6] David Cossock and Tong Zhang. 2008. Statistical Analysis of Bayes Optimal Subset Ranking. *IEEE Transactions on Information Theory* 54, 11 (2008), 5140–5154. <https://doi.org/10.1109/TIT.2008.929939>
- [7] Zhuyun Dai, Yubin Kim, and Jamie Callan. 2017. Learning To Rank Resources. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (Shinjuku, Tokyo, Japan) (SIGIR '17)*. Association for Computing Machinery, New York, NY, USA, 837–840. <https://doi.org/10.1145/3077136.3080657>
- [8] Clemens Damke and Eyke Hüllermeier. 2021. Ranking Structured Objects with Graph Neural Networks. *CoRR* abs/2104.08869 (2021). [arXiv:2104.08869](https://arxiv.org/abs/2104.08869) <https://arxiv.org/abs/2104.08869>
- [9] Thomas Demeester, Dolf Trieschnigg, Dong Nguyen, Djoerd Hiemstra, and Ke Zhou. 2015. FedWeb Greatest Hits: Presenting the New Test Collection for Federated Web Search. In *Proceedings of the 24th International Conference on World Wide Web (Florence, Italy) (WWW '15 Companion)*. Association for Computing Machinery, New York, NY, USA, 27–28. <https://doi.org/10.1145/2740908.2742755>
- [10] Ronak Desai, Qi Yang, Zonghuan Wu, Weiyi Meng, and Clement Yu. 2006. Identifying Redundant Search Engines in a Very Large Scale Metasearch Engine Context. In *Proceedings of the 8th Annual ACM International Workshop on Web Information and Data Management (Arlington, Virginia, USA) (WIDM '06)*. Association for Computing Machinery, New York, NY, USA, 51–58. <https://doi.org/10.1145/1183550.1183562>
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ArXiv* abs/1810.04805 (2019).
- [12] Yongquan Dong, Eduard C. Dragut, and Weiyi Meng. 2019. Normalization of Duplicate Records from Multiple Sources. *IEEE Transactions on Knowledge and Data Engineering* 31, 4 (2019), 769–782. <https://doi.org/10.1109/TKDE.2018.2844176>
- [13] Eduard Dragut, B. Gupta, Brian Beirne, Ali Neyestani, Badr Atassi, Clement Yu, and Weiyi Meng. 2014. Merging Query Results From Local Search Engines for Georeferenced Objects. *ACM Transactions on the Web* 8 (10 2014). <https://doi.org/10.1145/2656344>
- [14] Sergey Edunov, Alexei Baevski, and Michael Auli. 2019. Pre-trained language model representations for language generation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. Association for Computational Linguistics, Minneapolis, Minnesota, 4052–4059. <https://doi.org/10.18653/v1/N19-1409>
- [15] Yixing Fan, Jiafeng Guo, Yanyan Lan, Jun Xu, Chengxiang Zhai, and Xueqi Cheng. 2018. Modeling Diverse Relevance Patterns in Ad-Hoc Retrieval. In *The 41st International ACM SIGIR Conference on Research Development in Information Retrieval (Ann Arbor, MI, USA) (SIGIR '18)*. Association for Computing Machinery, New York, NY, USA, 375–384. <https://doi.org/10.1145/3209978.3209980>
- [16] Sara Nouri Golmaei and Xiao Luo. 2021. DeepNote-GNN: predicting hospital readmission using clinical notes and patient network. *Proceedings of the 12th ACM Conference on Bioinformatics, Computational Biology, and Health Informatics (2021)*.
- [17] Fatih Hafızoğlu. 2018. *Improving the efficiency of distributed information retrieval using hybrid index partitioning*. Master's thesis. Middle East Technical University.
- [18] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive Representation Learning on Large Graphs. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. <https://proceedings.neurips.cc/paper/2017/file/5dd9db5e033da9c6f5ba83c7a7e99-Paper.pdf>
- [19] Baoli Han, Ling Chen, and Xiaoxue Tian. 2018. Knowledge Based Collection Selection for Distributed Information Retrieval. *Inf. Process. Manage.* 54, 1 (jan 2018), 116–128. <https://doi.org/10.1016/j.ipm.2017.10.002>
- [20] Soheil Hassanipour, Haleh Ghaem, Morteza Arab-Zozani, Mohzhan Seif, Mohammad Fararouei, Elham Abdzadeh, Golnar Sabetian, and Shahram Paydar. 2019. Comparison of artificial neural network and logistic regression models for prediction of outcomes in trauma patients: A systematic review and meta-analysis. *Injury* 50, 2 (2019), 244–250. <https://doi.org/10.1016/j.injury.2019.01.007>
- [21] Yixuan He, Quan Gan, David Wipf, Gesine Reinert, Junchi Yan, and Mihai Cucuringu. 2022. GNNRank: Learning Global Rankings from Pairwise Comparisons via Directed Graph Neural Networks. *CoRR* abs/2202.00211 (2022). [arXiv:2202.00211](https://arxiv.org/abs/2202.00211) <https://arxiv.org/abs/2202.00211>
- [22] Dzung Hong, Luo Si, Paul Bracke, Michael Witt, and Tim Juchcinski. 2010. A Joint Probabilistic Classification Model for Resource Selection. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval (Geneva, Switzerland) (SIGIR '10)*. Association for Computing Machinery, New York, NY, USA, 98–105. <https://doi.org/10.1145/1835449.1835468>
- [23] Ziniu Hu, Dong Yuxiao, Kuansan Wang, and Yizhou Sun. 2020. Heterogeneous Graph Transformer.
- [24] Shan Jin and Man Lan. 2014. Simple May Be Best - A Simple and Effective Method for Federated Web Search via Search Engine Impact Factor Estimation. In *Proceedings of The Twenty-Third Text Retrieval Conference, TREC 2014, Gaithersburg, Maryland, USA, November 19-21, 2014 (NIST Special Publication, Vol. 500-308)*, Ellen M. Voorhees and Angela Ellis (Eds.). National Institute of Standards and Technology (NIST). http://trec.nist.gov/pubs/trec23/papers/pro-ECNU_federated.pdf
- [25] Armand Joulin, Edouard Grave, Piotr Bojanowski, and Tomas Mikolov. 2017. Bag of Tricks for Efficient Text Classification. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*. Association for Computational Linguistics, Valencia, Spain, 427–431. <https://aclanthology.org/E17-2068>
- [26] Yubin Kim. 2019. Robust Selective Search. *SIGIR Forum* 52 (2019), 170–171.
- [27] Diederik Kingma and Jimmy Ba. 2014. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations* (12 2014).
- [28] Anagha Kulkarni, Almer S. Tigelaar, Djoerd Hiemstra, and Jamie Callan. 2012. Shard Ranking and Cutoff Estimation for Topically Partitioned Collections. In *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (Maui, Hawaii, USA) (CIKM '12)*. Association for Computing Machinery, New York, NY, USA, 555–564. <https://doi.org/10.1145/2396761.2396833>
- [29] Quoc V. Le and Tomas Mikolov. 2014. Distributed Representations of Sentences and Documents. In *ICML*.
- [30] Juanhui Li, Yao Ma, Wei Zeng, Suqi Cheng, Jiliang Tang, Shuaiqiang Wang, and Dawei Yin. 2022. Graph Enhanced BERT for Query Understanding. *arXiv preprint arXiv:2204.06522* (2022).
- [31] Ping Li, Christopher J. C. Burges, and Qiang Wu. 2007. McRank: Learning to Rank Using Multiple Classification and Gradient Boosting. In *Proceedings of the 20th International Conference on Neural Information Processing Systems (Vancouver, British Columbia, Canada) (NIPS '07)*. Curran Associates Inc., Red Hook, NY, USA, 897–904.
- [32] Fang Liu, Shuang Liu, Clement Yu, Weiyi Meng, Ophir Frieder, and David Grossman. 2005. Database Selection in Intranet Mediators for Natural Language Queries. In *Proceedings of the 14th ACM International Conference on Information and Knowledge Management (Bremen, Germany) (CIKM '05)*. Association for Computing Machinery, New York, NY, USA, 229–230. <https://doi.org/10.1145/1099554.1099601>
- [33] King-Lup Liu, Weiyi Meng, Jing Qiu, Clement Yu, Vijay Raghavan, Zonghuan Wu, Yiyao Lu, Hai He, and Hongkun Zhao. 2007. AllInOneNews: Development and Evaluation of a Large-Scale News Metasearch Engine. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data (Beijing, China) (SIGMOD '07)*. Association for Computing Machinery, New York, NY, USA, 1017–1028. <https://doi.org/10.1145/1247480.1247601>
- [34] Yiyao Lu, Zonghuan Wu, Hongkun Zhao, Weiyi Meng, King-Lup Liu, Vijay Raghavan, and Clement Yu. 2007. MySearchView: A Customized Metasearch Engine Generator. In *Proceedings of the 2007 ACM SIGMOD International Conference on Management of Data (Beijing, China) (SIGMOD '07)*. Association for Computing Machinery, New York, NY, USA, 1113–1115. <https://doi.org/10.1145/1247480.1247625>

- [35] Xinyu Ma, Jiafeng Guo, Ruqing Zhang, Yixing Fan, Xiang Ji, and Xueqi Cheng. 2021. PROP: Pre-Training with Representative Words Prediction for Ad-Hoc Retrieval. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining* (Virtual Event, Israel) (WSDM '21). Association for Computing Machinery, New York, NY, USA, 283–291. <https://doi.org/10.1145/3437963.3441777>
- [36] Weiyi Meng and Clement Yu. 2010. *Advanced Metasearch Engine Technology*. Vol. 2. <https://doi.org/10.2200/S00307ED1V01Y201011DTM011>
- [37] Weiyi Meng and Clement T. Yu. 2011. *Search Engine Selection*. Springer International Publishing, Cham, 37–59. https://doi.org/10.1007/978-3-031-01843-5_3
- [38] Tomas Mikolov, Kai Chen, G.s Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *Proceedings of Workshop at ICLR 2013* (01 2013).
- [39] Dae Hoon Park and Yi Chang. 2019. Adversarial Sampling and Training for Semi-Supervised Information Retrieval. In *The World Wide Web Conference* (San Francisco, CA, USA) (WWW '19). Association for Computing Machinery, New York, NY, USA, 1443–1453. <https://doi.org/10.1145/3308558.3313416>
- [40] Abdel Naser Pouamoun and İlker Kocabaş. 0. Multi-agent-based hybrid peer-to-peer system for distributed information retrieval. *Journal of Information Science* 0, 0 (0), 01655515211010392. <https://doi.org/10.1177/01655515211010392> arXiv:<https://doi.org/10.1177/01655515211010392>
- [41] Priya Ranganathan, CS Pramesh, and Rakesh Aggarwal. 2017. Common pitfalls in statistical analysis: Logistic regression. *Perspectives in clinical research* 8, 3 (2017), 148–151. https://doi.org/10.4103/picr.picr_87_17
- [42] Aravind Sankar, Yozen Liu, Jun Yu, and Neil Shah. 2021. Graph Neural Networks for Friend Ranking in Large-Scale Social Platforms. In *Proceedings of the Web Conference 2021* (Ljubljana, Slovenia) (WWW '21). Association for Computing Machinery, New York, NY, USA, 2535–2546. <https://doi.org/10.1145/3442381.3450120>
- [43] Michael Schlichtkrull, Thomas N. Kipf, Peter Bloem, Rianne van den Berg, Ivan Titov, and Max Welling. 2018. Modeling Relational Data with Graph Convolutional Networks. In *The Semantic Web - 15th International Conference, ESWC 2018, Proceedings (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics))*. Springer/Verlag, 593–607. https://doi.org/10.1007/978-3-319-93417-4_38 15th International Conference on Extended Semantic Web Conference, ESWC 2018 ; Conference date: 03-06-2018 Through 07-06-2018.
- [44] Deanna Schreiber-Gregory and Karlen Bader. 2018. Logistic and Linear Regression Assumptions: Violation Recognition and Control.
- [45] Milad Shokouhi. 2007. Central-Rank-Based Collection Selection in Uncooperative Distributed Information Retrieval. *Proceedings of ECIR 4425*, 160–172. https://doi.org/10.1007/978-3-540-71496-5_17
- [46] Milad Shokouhi and Luo Si. 2011. Federated Search. *Foundations and Trends® in Information Retrieval* 5, 1 (2011), 1–102. <https://doi.org/10.1561/1500000010>
- [47] Luo Si and Jamie Callan. 2003. Relevant Document Distribution Estimation Method for Resource Selection. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval* (Toronto, Canada) (SIGIR '03). Association for Computing Machinery, New York, NY, USA, 298–305. <https://doi.org/10.1145/860435.860490>
- [48] Paul Thomas and Milad Shokouhi. 2009. SUSHI: Scoring Scaled Samples for Server Selection. 419–426. <https://doi.org/10.1145/1571941.1572014>
- [49] Günter Urak, Hermann Ziak, and Roman Kern. 2018. Source Selection of Long Tail Sources for Federated Search in an Uncooperative Setting. In *Proceedings of the 33rd Annual ACM Symposium on Applied Computing* (Pau, France) (SAC '18). Association for Computing Machinery, New York, NY, USA, 720–727. <https://doi.org/10.1145/3167132.3167212>
- [50] Tianfeng Wu, Xiaofeng Liu, and Shoubin Dong. 2019. LTRRS: A Learning to Rank Based Algorithm for Resource Selection in Distributed Information Retrieval. 52–63. https://doi.org/10.1007/978-3-030-31624-2_5
- [51] Jinxi Xu and W. Croft. 2000. Cluster-based Language Models for Distributed Retrieval. (02 2000). <https://doi.org/10.1145/312624.312687>
- [52] Clement Yu, George Philip, and Weiyi Meng. 2003. Distributed Top-N Query Processing with Possibly Uncooperative Local Systems. In *Proceedings of the 29th International Conference on Very Large Data Bases - Volume 29* (Berlin, Germany) (VLDB '03). VLDB Endowment, 117–128.
- [53] Xueli Yu, Weizhi Xu, Zeyu Cui, Shu Wu, and Liang Wang. 2021. Graph-based Hierarchical Relevance Matching Signals for Ad-hoc Retrieval. *CoRR* abs/2102.11127 (2021). arXiv:2102.11127 <https://arxiv.org/abs/2102.11127>
- [54] Jing Yuan, Lihong He, Eduard C. Dragut, Weiyi Meng, and Clement Yu. 2017. Result Merging for Structured Queries on the Deep Web with Active Relevance Weight Estimation. *Information Systems* 64 (2017), 93–103. <https://doi.org/10.1016/j.is.2016.06.005>