Entity Set Search of Scientific Literature: An Unsupervised Ranking Approach

Jiaming Shen, Jinfeng Xiao, Xinwei He, Jingbo Shang, Saurabh Sinha, Jiawei Han

Department of Computer Science, University of Illinois Urbana-Champaign, IL, USA

{js2, jxiao13, xhe17, shang7, sinhas, hanj}@illinois.edu

ABSTRACT

Literature search is critical for any scientific research. Different from Web or general domain search, a large portion of queries in scientific literature search are *entity-set queries*, that is, *multiple entities of possibly different types*. Entity-set queries reflect user's need for finding documents that contain multiple entities and reveal inter-entity relationships and thus pose non-trivial challenges to existing search algorithms that model each entity separately. However, entity-set queries are usually sparse (*i.e.*, not so repetitive), which makes ineffective many supervised ranking models that rely heavily on associated click history. To address these challenges, we introduce SetRank, an unsupervised ranking framework that models inter-entity relationships and captures entity type information. Furthermore, we develop a novel unsupervised model selection algorithm, based on the technique of weighted rank aggregation, to automatically choose the parameter settings in SetRank without resorting to a labeled validation set. We evaluate our proposed unsupervised approach using datasets from TREC Genomics Tracks and Semantic Scholar's query log. The experiments demonstrate that SetRank significantly outperforms the baseline unsupervised models, especially on entity-set queries, and our model selection algorithm effectively chooses suitable parameter settings.

KEYWORDS

Entity-Set Aware Search; Unsupervised Ranking Model; Unsupervised Model Selection; Literature Search

ACM Reference Format:

Jiaming Shen, Jinfeng Xiao, Xinwei He, Jingbo Shang, Saurabh Sinha, Jiawei Han. 2018. Entity Set Search of Scientific Literature: An Unsupervised Ranking Approach. In *SIGIR '18: The 41st International ACM SIGIR Conference on Research and Development in Information Retrieval*, July 8–12, 2018, Ann Arbor, MI, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3209978.3210055>

1 INTRODUCTION

Literature search helps a researcher identify relevant papers and summarize essential claims about a topic, forming a critical step in any scientific research. With the fast-growing volume of scientific publications, a good literature search engine is essential to

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

SIGIR '18, July 8–12, 2018, Ann Arbor, MI, USA

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-5657-2/18/07...\$15.00

<https://doi.org/10.1145/3209978.3210055>

Table 1: Ranking performance on 100 benchmark queries of the S2 production system. Entity-set queries (ESQs), marked bold, perform much weaker than non-ESQs do.

Metrics	ESQs	non-ESQs	Overall
NDCG@5	0.3622	0.6291	0.5223
NDCG@10	0.3653	0.6286	0.5233
NDCG@15	0.3840	0.6221	0.5269
NDCG@20	0.4011	0.6247	0.5353

researchers, especially in the domains like computer science and biomedical science where the literature collections are so massive, diverse, and rapidly evolving—few people can master the state-of-the-art comprehensively and in depth.

A large set of literature search queries contain multiple entities which can be either concrete instances (*e.g.*, *GABP* (a gene)) or abstract concepts (*e.g.*, *clustering*). We refer these queries as *entity-set queries*. For example, a computer scientist may want to find out how *knowledge base* can be used for *document retrieval* and thus issues a query “*knowledge base for document retrieval*”, which is an entity-set query containing two entities. Similarly, a biologist may want to survey how genes *GABP*, *TERT*, and *CD11b* are associated with *cancer* and submits a query “*GABP TERT CD11b cancer*”, another entity-set query with one disease and three gene entities.

Compared with typical short keyword queries, a distinctive characteristic of entity-set queries is that they reflect user's need for finding documents containing inter-entity relations. For example, among 50 queries collected from biologists in 2005 as part of TREC Genomics Track [14], 40 of them are explicitly formulated as finding relations among at least two entities. In most cases, a user who submits an entity-set query will expect to get a ranked list of documents that are most relevant to the whole *entity set*. Therefore, as in the previous examples, returning a paper about only *knowledge bases* or only one gene *GABP* is unsatisfactory.

Entity-set queries pose non-trivial challenges to existing search platforms. For example, among the 100 queries¹ released by Semantic Scholar (S2), 40 of them are entity-set queries and S2's production ranking system performs poorly on these entity-set queries, as shown in Table 1. The difficulties of handling entity-set queries mainly come from two aspects. First, entity relations within entity sets have not been modeled effectively. The association or co-occurrence of multiple entities has not gained adequate attention from existing ranking models. As a result, those models will rank papers where a single distinct entity appears multiple times higher than those containing many distinct entities. Second, entity-set queries are particularly challenging for supervised ranking models. As manual labeling of document relevance in academic search requires domain expertise, it is too expensive to train a ranking model

¹<http://data.allenai.org/esr/Queries/>

based purely on manually labeling. Most systems will first apply an off-the-shelf unsupervised ranking model during their *cold-start* process and then collect user interaction data (e.g., click information). Unfortunately, entity-set queries are usually sparse (i.e., not so repetitive), and have less associated click information. Furthermore, many off-the-shelf unsupervised models cannot return reasonably good candidate documents for entity-set queries within the top-20 positions. Many highly relevant documents will not be presented to users, which further compromises the usefulness of clicking information.

This paper tackles the new challenge—*improving the search quality of scientific literature on entity-set queries* and proposes an unsupervised ranking approach. We introduce SetRank, an unsupervised ranking framework that explicitly models inter-entity relations and captures entity type information. SetRank first links entity mentions in query and documents to an external knowledge-base. Then, each document is represented with both bag-of-words and bag-of-entities representations [36, 37] and fits two language models respectively. On the query side, a novel heterogeneous graph representation is proposed to model complex entity information (e.g., entity type) and entity relations within the set. This heterogeneous query graph represents all the information need in that query. Finally, the query-document matching is defined as a *graph covering process* and each document is ranked based on the information need it covers in the query graph.

Although being an unsupervised ranking framework, SetRank still has some parameters that need to be appropriately learned using a labeled validation set. To further automate the process of ranking model development, we develop a novel unsupervised model selection algorithm based on the technique of weighted rank aggregation. Given a set of queries with no labeled documents, and a set of candidate parameter settings, this algorithm automatically learns the most suitable parameter settings for that set of queries.

The significance of our proposed unsupervised ranking approach is two-fold. First, SetRank itself, as an unsupervised ranking model, boosts the literature search performance on entity-set queries. Second, SetRank can be adopted during the *cold-start* process of a search system, which enables the collection of high-quality click data for training subsequent supervised ranking model. Our experiments on two benchmark datasets² demonstrate the usefulness of our unsupervised model selection algorithm and the effectiveness of SetRank for searching scientific literature, especially on entity-set queries.

In summary, this work makes the following contributions:

- (1) A new research problem, effective entity-set search of scientific literature, is studied.
- (2) SetRank, an unsupervised ranking framework, is proposed, which models inter-entity relations and captures entity type information.
- (3) A novel unsupervised model selection algorithm is developed, which automatically selects SetRank's parameter settings without resorting to a labeled validation set.
- (4) Extensive experiments are conducted in two scientific domains, demonstrating the effectiveness of SetRank and our unsupervised model selection algorithm.

²Both benchmark datasets and our model implementations are publicly available at: <https://github.com/mickeystroller/SetRank>.

The remaining of the paper is organized as follows. Section 2 discusses related work. Section 3 presents our ranking framework SetRank. Section 4 presents the unsupervised model selection algorithm. Section 5 reports and analyzes the experimental results on two benchmark datasets and shows a case study of SetRank for biomedical literature search. Finally, Section 6 concludes this work with discussions on some future directions.

2 RELATED WORK

We examine related work in three aspects: academic search, entity-aware ranking model, and automatic ranking model selection.

2.1 Academic Search

The practical importance of finding highly relevant papers in scientific literature has motivated the development of many academic search systems. Google Scholar is arguably the most widely used system due to its large coverage. However, the ranking result of Google Scholar is still far from satisfactory because of its bias toward highly cited papers [1]. As a result, researchers may choose other academic search platforms, such as CiteSeerX [33], AMiner [30], PubMed [20], Microsoft Academic Search [29] and Semantic Scholar [38]. Research efforts of many such systems focus on the analytical tasks of scholar data such as author name disambiguation [30], paper importance modeling [28], and entity-based distinctive summarization [26]. However, this work focuses on ad-hoc document retrieval and ranking in academic search. The most relevant work to ours is [38] in which entity embeddings are used to obtain “soft match” feature of each (query, document) pair. However, [38] requires training data to combine word-based and entity-based relevance scores and to select parameter settings, which is rather different from our unsupervised approach.

2.2 Entity-aware Ranking Model

Entities, such as people, locations, or abstract concepts, are natural units for organizing and retrieving information [10]. Previous studies found that over 70% of Bing's query and more than 50% of traffic in Semantic Scholar are related to entities [12, 38]. The recent availability of large-scale knowledge repositories and accurate entity linking tools have further motivated a growing body of work on entity-aware ranking models. These models can be roughly categorized into three classes: expansion-based, projection-based, and representation-based.

The expansion-based methods use entity descriptions from knowledge repositories to enhance query representation. Xu *et al.* [39] use entity descriptions in Wikipedia as pseudo relevance feedback corpus to obtain cleaner expansion terms; Xiong and Callen [35] utilize the description of Freebase entities related to the query for query expansion; Dalton *et al.* [7] expand a query using the text fields of the attributes of the query-related entities and generate richer learning-to-rank features based on the expanded texts.

The projection-based methods try to project both query and document onto an entity space for comparison. Liu and Fang [19] use entities from a query and its related documents to construct a latent entity space and then connect the query and documents based on the descriptions of the latent entities. Xiong and Callen [34] use the textual features among query, entities, and documents to model the query-entity and entity-document connections. These additional

connections between query and document are then utilized in a learning-to-rank model. A fundamental difference of our work from the above methods is that we do not represent query and document using external terms/entities that they do not contain. This is to avoid adding noisy expansion of terms/entities that may not reflect the information need in the original user query.

The representation-based methods, as a recent trend for utilizing entity information, aim to build entity-enhanced text representation [37]. Xiong *et al.* [36] propose a bag-of-entities representation and demonstrated its effectiveness for vector space model. Raviv *et al.* [25] leverage the surface names of entities to build an entity-based language model. Many supervised ranking models are proposed to apply learning-to-rank methods for combining entity-based signals with word-based signals. For example, ESR [38] uses entity embeddings to compute entity-based query-document matching score and then combines it with word-based score using RankSVM. Following the same spirit, Xiong *et al.* [37] propose a word-entity duet framework that simultaneously models the entity annotation uncertainty and trains the ranking model. Comparing with the above methods, we also use the bag-of-entity representation but combine it with word-based representation in an unsupervised way. Also, to the best of our knowledge, we are the first to capture entity relation and type information in an unsupervised entity-aware ranking model.

2.3 Automatic Ranking Model Selection

Most ranking models need to manually set many parameter values. To automate the process of selecting parameter settings, some AutoML methods [3, 8] are proposed. Nevertheless, these methods still require a validation set which contains queries with labeled documents. In this paper, we develop an unsupervised model selection algorithm, based on rank aggregation, to automatically choose parameter settings without resorting to a labeled validation set. Rank aggregation aims to combine multiple existing rankings into a joint ranking. Fox and Shaw [9] propose some deterministic functions to combine rankings heuristically. Klementiev *et al.* [17, 18] propose an unsupervised learning algorithm for rank aggregation based on a linear combination of ranking functions. Another related line of work is to model rankings using a statistical model (e.g., Plackett-Luce model) and aggregate them based on statistical inference [11, 21, 41]. Lately, Bhowmik and Ghosh [2] propose to use object attributes to augment some standard rank aggregation framework. Compared with these methods, our proposed algorithm goes beyond just combining multiple rankings and uses aggregated ranking to guide the selection of parameter settings.

3 RANKING FRAMEWORK

This section presents our unsupervised ranking framework for leveraging entity (set) information in search. Our framework provides a principled way to rank a set of documents D for a query q . In this framework, we represent each document using standard bag-of-words and bag-of-entities representations [36, 37] (Section 3.1) and represent the query using a novel heterogeneous graph (Section 3.2) which naturally model the entity set information. Finally, we model the query-document matching as a “graph covering” process, as described in Section 3.3.

Figure 1: An illustrative example showing one document comprised of two fields (i.e., title, abstract) with their corresponding bag-of-words and bag-of-entities representations.

3.1 Document Representation

We represent each document using both word and entity information. For words, we use standard bag-of-words representation and treat each unigram as a word. For entities, we adopt an entity linking tool (details described in Section 5.2) that utilizes a knowledge base/graph (e.g., Wikidata or Freebase) where entities have unique IDs. Given an input text, this tool will find the entity mentions (i.e., entity surface names) in the text and link each of them to a disambiguated entity in the knowledge base/graph. For example, given the input document title “*Training linear SVMs in linear time*”, this tool will link the entity mention “SVMs” to the entity “Support Vector Machine” with Freebase id ‘/m/0hc2f’. Previous studies [25, 36] show that when the entity linking error is within a reasonable range, the returned entity annotations, though noisy, can still improve the overall search performance, partially due to the following:

- (1) **Polysemy resolution.** Different entities with the same surface name will be resolved by the entity linker. For example, the fruit “Apple” (with id ‘/m/014j1m’) will be disambiguated with the company “Apple” (with id ‘/m/0k8z’).
- (2) **Synonymy resolution.** Different entity surface names corresponding to the same entity will be identified and merged. For example, the entity “United States of America” (with id ‘/m/09c7w0’) can have different surface names including “USA”, “United States”, and “U.S.” [25]. The entity linker can map all these surface names to the same entity.

After linking all the entity mentions in a document to entities in the knowledge base, we can obtain the bag-of-entities representation of this document. Then, we fit two language models (LMs) for this document: one being word-based (i.e., traditional unigram LM) and the other being entity-based. Notice that in the literature search scenario, documents (i.e., papers) usually contain multiple fields, such as title, abstract, and full text. We model each document field using a separate bag-of-words representation and a separate bag-of-entities representation, as shown in Figure 1.

To exploit such intra-document structures, we generally assume a document d_i has k fields $d_i = \{d_{i,1}, \dots, d_{i,k}\}$ and thus the document collection can be separated into k parts: $\{D_1, \dots, D_k\}$. Following [23], we assign each field a weight δ_j and formulate the generation process of a token t given the document d_i as follows:

$$p(t|d_i) = \sum_{j=1}^k p(t|d_{i,j})p(d_{i,j}|d_i), \quad p(d_{i,j}|d_i) = \frac{\delta_j}{\sum_{j'=1}^k \delta_{j'}}. \quad (1)$$

Notice the a token t can be either a unigram w or an entity e , and the field weight δ_j can be either manually set based on prior knowledge or automatically learned using the mechanism described in Section 4. The token generation probability under each document field $p(t|d_{i,j})$ can be obtained from the maximum likelihood estimate with Dirichlet prior smoothing [40] as follows:

$$p(t|d_{i,j}) = \frac{n_{t,d_{i,j}} + \mu_j \frac{n_{t,D_j}}{L_{D_j}}}{L_{d_{i,j}} + \mu_j}, \quad (2)$$

where $n_{t,d_{i,j}}$ and $L_{d_{i,j}}$ represent the number of token t in $d_{i,j}$ and the length of $d_{i,j}$. Similarly, we can define n_{t,D_j} and L_{D_j} . Finally, μ_j is a scale parameter of the Dirichlet distribution for field j . A concrete example is shown in Figure 1.

3.2 Query Representation

Given an input query q , we first apply the same entity linker used for document representation to extract all the entity information in the query. Then, we design a novel heterogeneous graph to represent this query q , denoted as G_q . Such a graph representation captures both word and entity information in the query and models the entity relations. A concrete example is shown in Figure 2.

Node representation. In this heterogeneous query graph, each node represents a query token. As a token can be either a word or an entity, there are two different types of nodes in this graph.

Edge representation. We use an edge to represent a latent relation between two query tokens. In this work, we consider two types of latent relations: word-word relation and entity-entity relation. For word-word relation, we add an edge for each pair of adjacent word tokens with equal weight 1. For instance, given an query “*Atari video games*”, we will add two edges, one between word pairs $\langle \text{Atari}, \text{video} \rangle$ and the other between $\langle \text{video}, \text{game} \rangle$. On the entity side, we aim to emphasize all the possible entity-entity relations, and thus add an edge between each pair of entity tokens.

Modeling entity type. The type information of each query entity can further reveal the user’s information need. Therefore, we assign the weight of each entity-entity relation based on these two entities’ type information. Intuitively, if the types of two entities are distant from each other in a type hierarchy, then the relation between these two entities should have a larger weight. A similar idea is exploited in [10] and found useful for type-aware entity retrieval.

Mathematically, we use ϕ_e to denote the type of entity e ; use $LCA_{u,v}$ to denote the Lowest Common Ancestor (LCA) of two nodes u and v in a given tree (i.e., type hierarchy), and use $l(u,v)$ to denote the length of a path between node u and node v . In Figure 2, for example, the entity tokens ‘/m/0hjlw’ and ‘/m/0xwj’, corresponding to “*reinforcement learning*” and “*Atari*”, have types ‘education.field_of_study’ and ‘computer.game’, respectively. The Lowest Common Ancestor of these two types in the type hierarchy is ‘Thing’. Finally, we define the relation strength between entity

Figure 2: An illustrative example showing the heterogeneous graph representation of one query. Word-word relations are marked by dash lines and entity-entity relations are marked by solid lines. Different solid line colors represent different relation strengths based on two entities’ types.

e_1 and entity e_2 as follows:

$$LCA_{e_1, e_2} = LCA(\phi_{e_1}, \phi_{e_2}), \quad (3)$$

$$\lambda_{e_1, e_2} = 1 + \max \{l(\phi_{e_1}, LCA_{e_1, e_2}), l(\phi_{e_2}, LCA_{e_1, e_2})\}. \quad (4)$$

Our proposed heterogeneous query graph representation is general and can be extended. For example, we can apply dependency parsing for verbose queries, and only add an edge between two word tokens that have direct dependency relation. Also, if the importance of each entity-entity relation is given, we can then set the edge weights correspondingly. We leave these extensions for future works.

3.3 Document Ranking using Query Graph

Our proposed heterogeneous query graph G_q represents all information need in the user-issued query. Such need can be either to find document discussing one particular entity or to identify papers studying an important inter-entity relation. Intuitively, a document that can satisfy more information need should be ranked at a higher position. To quantify such information need that is explained by a document, we define the following graph covering process.

Query graph covering. If a query token $t \in q$ exists in a document d_i , we say d_i covers the node in G_q that corresponds to this token. Similarly, if a pair of query tokens t_1 and t_2 exists in d_i , we say d_i covers the edge in G_q that corresponds to the relation of this token pair $\langle t_1, t_2 \rangle$. The subgraph of G_q that is covered by the document d_i , denoted as $G_{q|d_i}$, represents the information need in the query q that is explained by the document d_i .

Furthermore, we follow the same spirit of [22] and view the subgraph $G_{q|d_i}$ as a Markov Network, based on which we define the joint probability of the document d_i and the query q as follows:

$$P(d_i, q) \stackrel{\text{def}}{=} \frac{1}{Z} \prod_{c \in G_{q|d_i}} \psi(c) \stackrel{\text{rank}}{=} \sum_{c \in G_{q|d_i}} \log \psi(c) \stackrel{\text{rank}}{=} \sum_{c \in G_{q|d_i}} f(c), \quad (5)$$

where Z is a normalization factor, c indexes the cliques in graph, and $\psi(c)$ is the non-negative potential defined on c . The last equation holds as we let $\psi(c) = \exp[f(c)]$. Notice that if $G_{q|d_1}$ is a subgraph of $G_{q|d_2}$ which means document d_1 covers less information than document d_2 does, we should have $P(d_1, q) < P(d_2, q)$. Therefore, we should design $f(\cdot)$ to satisfy the constraint $f(c) > 0, \forall c$.

In this work, we focus on modeling each single entity and pairwise relations between two entities. Therefore, each clique c can be either a node or an edge in the graph. Modeling higher-order

relations among more than two entities (*i.e.*, cliques with size larger than 2) is left for future work. We define the potential functions for a single node and an edge as follows:

Node potential. Node potential quantifies the information need contained in a single node t , which can be either a word token w or an entity token e . To balance the relative weight of a word token and an entity token, we introduce a parameter $\lambda_E \in [0, 1]$, and define the node potential function $f(\cdot)$ as follows:

$$f(t) = \begin{cases} \lambda_E \cdot a(P(t|d_i)) & \text{if token } t \text{ is an entity token} \\ (1 - \lambda_E) \cdot a(P(t|d_i)) & \text{if token } t \text{ is a word token} \end{cases} \quad (6)$$

where $a(\cdot)$ is an activation function that transforms a raw probability to a node potential. Here, we set $a(x) = \sqrt{x}$ in order to amplify $P(t|d_i)$ which has a relatively small value.

Edge potential. Edge potential quantifies the information need contained in an edge $\langle t_1, t_2 \rangle$ that can be either a word-word (W-W) relation or an entity-entity (E-E) relation. In our query graph representation, all word-word relations have an equal weight of 1, and the weight of each entity-entity relation (*i.e.*, λ_{e_1, e_2}) is defined by Equation (3). Finally, we calculate the edge potential as follows:

$$f(\langle t_1, t_2 \rangle) = \lambda_{\langle t_1, t_2 \rangle} \cdot a(P(t_1, t_2|d_i)), \quad (7)$$

$$\lambda_{\langle t_1, t_2 \rangle} = \begin{cases} \lambda_E \cdot \lambda_{e_1, e_2} & \text{if } \langle t_1, t_2 \rangle \text{ is an E-E relation} \\ (1 - \lambda_E) & \text{if } \langle t_1, t_2 \rangle \text{ is a W-W relation} \end{cases} \quad (8)$$

where $\lambda_{\langle t_1, t_2 \rangle}$ measures the edge importance, and $a(\cdot)$ is the same activation function as defined above. To simplify the calculation of $P(t_1, t_2|d_i)$, we make an assumption that two tokens t_1 and t_2 are conditionally independent given a document d_i . Then, we replace $P(t_1, t_2|d_i)$ with $P(t_1|d_i)P(t_2|d_i)$ and substitute it in Equation (7).

Putting all together. After defining the node and edge potentials, we can calculate the joint probability of each document d_i and query q using Equation (5) as follows:

$$P(d_i, q) = (1 - \lambda_E) \sum_{w \in G_q|d_i} \left(1 + \sum_{\langle w, w' \rangle \in G_q|d_i} a(P(w'|d_i)) \right) a(P(w|d_i)) \\ + \lambda_E \sum_{e \in G_q|d_i} \left(1 + \sum_{\langle e, e' \rangle \in G_q|d_i} \lambda_{e, e'} \cdot a(P(e'|d_i)) \right) a(P(e|d_i)). \quad (9)$$

As shown in the above equation, SetRank will explicitly reward paper capturing inter-entity relations and covering more unique entities. Also, it uses λ_E to balance the word-based relevance with entity-based relevance, and models entity type information in $\lambda_{e, e'}$.

4 UNSUPERVISED MODEL SELECTION

Although being an unsupervised ranking framework, SetRank still has some parameters that need to be appropriately set by ranking model designers, including the weight of title/abstract field and the relative importance of entity token λ_E . Previous study [40] shows that these model parameters have significant influences on the ranking performance and thus we need to choose them carefully. Typically, these parameters are chosen to optimize the performance over a validation set that is manually constructed and contains the relevance label of each query-document pair. Though being useful, the validation set is not always available, especially for those applications (*e.g.*, literature search) where labeling document requires domain expertise.

To address the above problem, we propose an unsupervised model selection algorithm which automatically chooses the parameter settings without resorting to a manually labeled validation set. The key philosophy is that although people who design the ranking model do not know the exact “optimal” parameter settings, they do have prior knowledge about the reasonable range for each of them. For example, the title field weight should be set larger than the abstract field weight, and the entity token weight λ_E should be set small if the returned entity linking results are noisy. Our model selection algorithm leverages such prior knowledge by letting the ranking model designer input the search range of each parameter’s value. It will then return the best value for each parameter within its corresponding search range. We first describe our notations and formulate our problem in Section 4.1. Then, we present our model selection algorithm in Section 4.2.

4.1 Notations and Problem Formulation

Notations. We use S_K to denote the collection of rankings over a set of K documents: $D = \{d_1, \dots, d_k, \dots, d_K\}, k \in [K] = \{1, \dots, K\}$. We denote by $\pi : [K] \rightarrow [K]$ a *complete ranking*, where $\pi(k)$ denotes the position of document d_k in the ranking, and $\pi^{-1}(j)$ is the index of the document on position j . For example, given the ranking: $d_3 > d_1 > d_2 > d_4$, we will have $\pi = [2, 3, 1, 4]$ and $\pi^{-1} = (3, 1, 2, 4)$. Furthermore, we use the symbol τ (instead of π) to denote an *incomplete ranking* which includes only some of the documents in D . If document d_k does not occur in the ranking, we set $\tau(k) = 0$, otherwise, $\tau(k)$ is the rank of document d_k . In the corresponding τ^{-1} , those missing documents simply do not occur. For example, given the ranking: $d_4 > d_2 > d_1$, we have $\tau = [3, 2, 0, 1]$ and $\tau^{-1} = (4, 2, 1)$. Finally, we let $I(\tau)$ to represent the index of documents that appear in the ranking list τ .

Problem Formulation. Given a parameterized ranking model M_θ where θ denotes the set of all parameters (*e.g.*, $\{k, b\}$ in BM25, $\{\mu\}$ in query likelihood model with dirichlet prior smoothing), we want to find the best parameter settings θ^* such that the ranking model M_{θ^*} achieves the best ranking performance over the space $\mathbb{Q}$ of all queries. In practice, however, the space consisting of all possible values of θ can be infinite and we cannot access all queries in $\mathbb{Q}$. Therefore, we assume ranking model designers will input p possible sets of parameter values: $\Theta = \{\theta_1, \dots, \theta_p\}$ and a finite subset of queries $Q \subset \mathbb{Q}$. Finally, we formulate our problem of *unsupervised model selection* as follows:

DEFINITION 1. (PROBLEM FORMULATION). Given a parameterized ranking model M_θ , p candidate parameter settings Θ , and an unlabeled query subset Q , we aim to find $\theta^* \in \Theta$ such that M_{θ^*} achieves the best ranking performance over Q .

4.2 Model Selection Algorithm

Our framework measures the goodness of each parameter settings $\theta_i \in \Theta$ based on its induced ranking model M_{θ_i} . The key challenge here is how we can evaluate the ranking performance of each M_{θ_i} over a query q which has no labeled documents. To address this challenge, we first leverage a weighted rank aggregation technique to obtain an aggregated rank list and then evaluate the quality of each M_{θ_i} based on the agreement between its generated rank list and the aggregated rank list. The key intuition here is that high-quality ranking models will rank documents based on a similar

Algorithm 1: Unsupervised Model Selection.

Input: A parameterized ranking model M_θ , p candidate parameter settings $\Theta = \{\theta_1, \dots, \theta_p\}$, and an unlabeled query subset Q .

Output: The best ranking model M_{θ^*} with $\theta^* \in \Theta$.

```

1 set  $score(M_{\theta_1}) = score(M_{\theta_2}) = \dots = score(M_{\theta_p}) = 0$ ;
2 for query  $q \in Q$  do
3   set  $\alpha_1 = \alpha_2 = \dots = \alpha_p = \frac{1}{p}$ ;
4   set  $\pi^{prev} = None$ ;
5   while True do
6     // Weighted Rank Aggregation ;
7     for document index  $j$  from 1 to  $|D|$  do
8        $score(d_j) = 0$ ;
9       for rank list index  $i$  from 1 to  $p$  do
10        if  $j \in I(\tau_i)$  (i.e.,  $d_j$  appears in  $\tau_i$ ) then
11           $score(d_j) = score(d_j) + \alpha_i(|\tau_i| + 1 - \tau_i(d_j))$ ;
12       $\pi = \text{argsort}(score(d_1), \dots, score(d_{|D|}))$ ;
13      // Confidence Score Adjustment ;
14      for rank list index  $i$  from 1 to  $p$  do
15         $\alpha_i = \frac{\exp(-dist(\tau_i || \pi))}{\sum_{i'} \exp(-dist(\tau_{i'} || \pi))}$ ;
16      // Convergence Check ;
17      if  $\pi == \pi^{prev}$  then
18        Break;
19      else
20         $\pi^{prev} \leftarrow \pi$ ;
21      for rank list index  $i$  from 1 to  $p$  do
22         $score(M_{\theta_i}) = score(M_{\theta_i}) + \alpha_i$ ;
23   $M_{\theta^*} = \arg \max_{\theta \in \Theta} score(M_\theta)$ ;
24  Return  $M_{\theta^*}$ ;
```

distribution while low-quality ranking models will rank documents in a uniformly random fashion. Therefore, the agreement between each rank list with the aggregated rank list serves as a good signal of its quality.

At a high level, our model selection method is an iterative algorithm which repeatedly aggregates multiple rankings (with their corresponding weights) and uses the aggregated rank list to estimate the quality of each of them. Given a query q , we first construct p ranking models M_{θ_i} , $i \in [1, \dots, p]$, one for each parameter settings $\theta_i \in \Theta$ and obtain its returned top- k rank list τ_i over a document set D_i (i.e., $|D_i| = k$). Then, we construct a unified document pool $D = \bigcup_{i=1}^p D_i$. After that, we use α_i to denote the confidence score of each ranking model M_{θ_i} , and initialize all of them with equal value $\frac{1}{p}$. During each iteration, we first aggregate $\{\tau_1, \dots, \tau_p\}$, weighted by $\{\alpha_1, \dots, \alpha_p\}$, and obtain the aggregated rank list π . Then, we adjust the confidence score of each ranking model M_{θ_i} (i.e., α_i) based on the distance of two rankings: τ_i and π . Here, we use π to denote the aggregated rank list because it is a complete ranking over the document pool D .

Weighted Rank Aggregation. We aggregate multiple rank lists using a variant of Borda counting method [6] which considers the relative weight of each rank list. We calculate the score of each document based on its position in each rank list as follows:

$$score(d_j) = \sum_{i=1}^p \alpha_i (|\tau_i| + 1 - \tau_i(d_j)) \mathbf{1}\{j \in I(\tau_i)\}, \quad (10)$$

where $|\tau_i|$ denotes the length of a rank list τ_i , and $\mathbf{1}\{x\}$ is an indicator function. When document d_j appears in the rank list τ_i , $\mathbf{1}\{j \in I(\tau_i)\}$

Figure 3: An illustrative example showing the process of weighted rank aggregation and the calculation of two different ranking distances (i.e., KT and $posKT$).

equals to 1, otherwise, it equals to 0. The above equation will reward a document ranked at higher position (i.e., small $\tau_i(d_j)$) in a high-quality rank list (i.e., large α_i) a larger score. Finally, we obtain the aggregated rank of these documents based on their corresponding scores. A concrete example is shown in Figure 3.

Confidence Score Adjustment. After we obtain the aggregated rank list, we will need to adjust the confidence score α_i of each ranking model M_{θ_i} based on the distance between τ_i and aggregated rank list π . In order to compare the distance between an incomplete rank list τ_i with a complete rank list π , we extend the classical Kendall Tau distance [16] and define it as follows:

$$KT(\tau_i || \pi) = \sum_{\substack{\tau_i(a) < \tau_i(b) \\ a, b \in I(\tau_i)}} \mathbf{1}\{\pi(a) > \pi(b)\}. \quad (11)$$

The above distance counts the number of pairwise disagreements between τ_i and π . One limitation of this distance is that it does not differentiate the importance of different ranking positions. Usually, switching two documents in the top part of a rank list should be penalized more, compared with switching another two documents in the bottom part of a rank list. To model such intuition, we propose a *position-aware* Kendall Tau distance and define it as follows:

$$posKT(\tau_i || \pi) = \sum_{\substack{\tau_i(a) < \tau_i(b) \\ a, b \in I(\tau_i)}} \left(\frac{1}{\log_2(1 + \pi(b))} - \frac{1}{\log_2(1 + \pi(a))} \right) \mathbf{1}\{\pi(a) > \pi(b)\}. \quad (12)$$

With the distance between two rankings defined, we can adjust the confidence score as follows:

$$\alpha_i = \frac{\exp(-dist(\tau_i || \pi))}{\sum_{i'} \exp(-dist(\tau_{i'} || \pi))}, \quad (13)$$

where $dist(\tau_i || \pi)$ can be either $KT(\tau_i || \pi)$ or $posKT(\tau_i || \pi)$ and we will study how different this choice can influence the model selection results in Section 5.4. The key idea of the above equation is to promote the ranking model which returns a ranked list better aligned with the aggregated rank list.

Putting all together. Algorithm 1 summarizes our unsupervised model selection process. Given a query $q \in Q$, we can iteratively apply weighted rank aggregation and confidence score adjustment until the algorithm converges. Then, we collect the converged $\{\hat{\alpha}_1, \dots, \hat{\alpha}_p\}$. Specifically, $\hat{\alpha}_i$ is the confidence score of ranking model M_{θ_i} on query q . With a slight abuse of notation, we use $score(M_{\theta_i})$ to denote its accumulated confidence score. Given a set of queries Q , we run the former procedure for each query and sum over all converged $\hat{\alpha}_i$. Finally, we return the ranking model M_{θ^*} which has the largest accumulated confidence score.

5 EXPERIMENTS

In this section, we evaluate our proposed SetRank framework as well as unsupervised model selection algorithm on two datasets from two scientific domains.

5.1 Datasets

We use two benchmark datasets for the experiments: Semantic Scholar [38] in Computer Science (S2-CS) and TREC 2004&2005 Genomics Track in Biomedical science (TREC-BIO).

S2-CS contains 100 queries sampled from Semantic Scholar’s query log, in which 40 queries are entity-set queries and the maximum number of entities in a query is 5. Candidate documents are generated by pooling from variations of Semantic Scholar’s online production system and all of them are manually labeled on a 5-level scale. Entities in both queries and documents are linked to Freebase using CMNS [13]. As the original dataset does not contain the entity type information, we enhance it by retrieving each entity’s most notable type in the latest Freebase dump³ based on its Freebase ID. These types are organized by Freebase type hierarchy.

TREC-BIO includes 100 queries designed by biologists and the candidate document pool is constructed based on the top results of all submissions at that time. All candidate documents are labeled on a 3-level scale. In these 100 queries, 86 of them are entity-set queries and the maximum number of entities in a query is 11. The original dataset contains no entity information and therefore we apply PubTator [32], the state-of-the-art biomedical entity linking tool, to obtain 5 types of entities (*i.e.*, *Gene*, *Disease*, *Chemical*, *Mutation*, and *Species*) in both queries and documents. We build a simple type hierarchy with root node named ‘Thing’ and each first-level node corresponds to one of the above 5 types.

5.2 Entity Linking Performance

We evaluate the query entity linking using precision and recall at the query level. Specifically, an entity annotation is considered correct if it appears in the gold labeled data (*i.e.*, the strict evaluation in [4]). The original S2-CS dataset provides such gold labeled data. For TREC-BIO dataset, we asked two Master-level students with biomedical science background to label all the linked entities as well as the entities that they could identify in the queries. We also report the entity linking performance on the general domain queries (ClueWeb09 and ClueWeb12) for references [36]. As we can see in Table 2, the overall linking performance of academic queries is better than that of general domain queries, probably because academic queries have less ambiguity. Also, recall of entity linking in TREC-BIO dataset is very high. A possible reason is that the biomedical entities have very distinctive tokens (*e.g.*, “*narcolepsy*” is a specific disease related to sleep and is seldom used in other contexts) and thus it is relatively easier to recognize them.

5.3 Ranking Performance

5.3.1 Experimental Setup.

Evaluation metrics. Since documents in both datasets have multi-level graded relevance, we use $NDCG@5,10,15,20$ as our main evaluation metrics. All evaluation is performed using standard

Table 2: Entity linking performance on scientific domain queries (S2-CS, TREC-BIO) and general domain queries (ClueWeb09, ClueWeb12).

	S2-CS	TREC-BIO	ClueWeb09	ClueWeb12
Precision	0.680	0.678	0.577	0.485
Recall	0.680	0.727	0.596	0.575

pytrec_eval tool [31]. Statistical significances are tested using two-tailed *t*-test with $p\text{-value} \leq 0.05$.

Baselines. We compare SetRank with 4 baseline ranking models: Vector Space Model (BM25 [27]), Query Likelihood Model with Dirichlet Prior smoothing (LM-DIR) or with Jelinek Mercer smoothing (LM-JM) [40], and the Information-Based model (IB) [5]. All models are applied to the paper’s title and abstract fields. Here, we do not compete with Semantic Scholar’s production system and ESR model [38] because they are supervised models trained over user’s click information which is not available in our setting.

The parameters of all models, including the field weights, are set using 5-fold cross validation over the queries in each benchmark dataset using the same paradigm in [25] as follows. For each hold-out fold, the other four folds are served as a validation set. A grid search is applied to choose the optimal parameter settings that maximize NDCG@20 on the validation set. Specifically, the title and abstract field weights are selected from {1,5,10,15,20,50}; the Dirichlet smoothing parameter μ and Jelinek Mercer smoothing parameter λ are chosen from {500, 1000, 1500, 2000, 2500, 3000} and {0.1, 0.2, ..., 0.9}, respectively; the relative weight of entity token λ_E used in SetRank is selected from {0, 0.1, ..., 1}. The best performing parameter settings are then saved for the hold-out evaluation.

5.3.2 Effectiveness of Leveraging Entity Information.

As mentioned before, the entity linking process is not perfect and it generates some noisy entity annotations. Therefore, we first study how different ranking models, including our proposed SetRank, can leverage such noisy entity information to improve the ranking performance. We evaluate three variations of each model – one using only word information, one using only entity information, and one using both pieces of information.

Results are shown in Table 3. We notice that the usefulness of entity information is *inconclusive* for baseline models. On S2-CS dataset, adding entity information can improve the ranking performance, while on TREC-BIO dataset, it will drag down the performance of all baseline methods. This resonates with previous findings in [15] that simply adding entities into queries and posting them to existing ranking models does not work for biomedical literature retrieval. Compared with baseline methods, SetRank successfully combines the word and entity information and effectively leverages such noisy entity information to improve the ranking performance. Furthermore, SetRank can better utilize each single information source, either word or entity, than other baseline models thanks to our proposed query graph representation. Overall, SetRank significantly outperforms all variations of baseline models.

5.3.3 Ranking Performance on Entity-Set Queries.

We further study each model’s ranking performance on entity-set queries. There are 40 and 86 entity-set queries in S2-CS and TREC-BIO, respectively. We denote these subsets of entity-set queries as S2-CS-ESQ and TREC-BIO-ESQ. As shown in Table 4, SetRank significantly outperforms the best variation of all baseline methods

³<https://developers.google.com/freebase/>

Table 3: Effectiveness of leveraging (noisy) entity information for ranking. Each method contains three variations and the best variation is labeled bold. The superscript “*” means the model significantly outperforms the best variation of all 4 baseline methods (with p-value ≤ 0.05).

Dataset	Method	BM25			LM-Dir			LM-JM			IB			SetRank		
		Word	Entity	Both	Word	Entity	Both	Word	Entity	Both	Word	Entity	Both	Word	Entity	Both
S2-CS	NDCG@5	0.3476	0.3319	0.3675	0.3447	0.3460	0.3563	0.3626	0.3394	0.3625	0.3759	0.3420	0.3729	0.3890	0.3761	0.4207*
	NDCG@10	0.3785	0.3520	0.4039	0.3623	0.3579	0.3901	0.3774	0.3519	0.3962	0.3903	0.3557	0.4009	0.4168	0.3885	0.4431*
	NDCG@15	0.4001	0.3616	0.4160	0.3781	0.3673	0.4077	0.4051	0.3666	0.4174	0.4113	0.3699	0.4272	0.4411	0.4054	0.4762*
	NDCG@20	0.4126	0.3752	0.4333	0.4012	0.3816	0.4205	0.4182	0.3804	0.4362	0.4295	0.3855	0.4421	0.4674	0.4229	0.4950*
TREC-BIO	NDCG@5	0.3189	0.1542	0.2613	0.3053	0.1755	0.2669	0.2957	0.1656	0.2826	0.3045	0.1842	0.2770	0.3417	0.2111	0.3744*
	NDCG@10	0.2968	0.1488	0.2472	0.2958	0.1601	0.2571	0.2742	0.1588	0.2572	0.2918	0.1715	0.2633	0.3165	0.1976	0.3522*
	NDCG@15	0.2833	0.1424	0.2395	0.2852	0.1579	0.2591	0.2642	0.1575	0.2437	0.2835	0.1664	0.2541	0.3017	0.1931	0.3363*
	NDCG@20	0.2739	0.1419	0.2337	0.2781	0.1558	0.2547	0.2560	0.1534	0.2362	0.2722	0.1628	0.2406	0.2900	0.1885	0.3246*

Table 4: Ranking performance on entity-set queries. The best variation of each baseline method is selected. The superscript “*” means the model significantly outperforms all 4 baseline methods (with p-value ≤ 0.05).

Dataset	Metric	BM25	LM-Dir	LM-JM	IB	SetRank
S2-CS -ESQ	NDCG@5	0.3994	0.3522	0.3812	0.3956	0.4983*
	NDCG@10	0.4364	0.3973	0.4241	0.4209	0.5130*
	NDCG@15	0.4454	0.4160	0.4431	0.4496	0.5450*
	NDCG@20	0.4609	0.4264	0.4618	0.4664	0.5629*
TREC-BIO -ESQ	NDCG@5	0.3185	0.2934	0.2940	0.3011	0.3639*
	NDCG@10	0.2968	0.2834	0.2746	0.2896	0.3406*
	NDCG@15	0.2812	0.2711	0.2636	0.2832	0.3251*
	NDCG@20	0.2718	0.2644	0.2553	0.2708	0.3132*

on S2-CS-ESQ and TREC-BIO-ESQ by at least 25% and 14% respectively in terms of NDCG@5. Also, we can see the advantages of SetRank over the baselines on entity-set queries are larger than those on general queries. This further demonstrates SetRank’s effectiveness of modeling entity set information.

5.3.4 Effectiveness of Modeling Entity Relation and Entity Type.

To study how the inter-entity relation and entity type information can contribute to document ranking, we compare SetRank with two of its variants, SetRank^{-t} and SetRank^{-ts}. The first variant models entity relation among the set but ignores the entity type information, and the second variant simply neglects both entity relation and type.

Results are shown in Table 5. First, we compare SetRank^{-t} with SetRank^{-ts} and find that modeling the entity relation in entity sets can significantly improve the ranking results. Such improvement is especially obvious on the entity-set query sets S2-CS-ESQ and TREC-BIO-ESQ. Also, by comparing SetRank with SetRank^{-t}, we can see adding entity type information can further improve ranking performance. In addition, we present a concrete case study for one entity-set query in Table 6. The top-2 papers returned by SetRank^{-ts} are focusing on *video game* without discussing its relation with *reinforcement learning*. In comparison, SetRank considers the entity relations and returns the paper mentioning both entities.

5.3.5 Analysis of Entity Token Weight λ_E .

We introduce the entity token weight λ_E in Eq. (6) to combine the entity-based and word-based relevance scores. In all previous experiments, we choose its value using cross validation. Here, we study how this parameter will influence the ranking performance by constructing multiple SetRank models with different λ_E and directly report their performance on all 100 queries.

As shown in Figure 4, for S2-CS dataset, SetRank’s ranking performance first increases as λ_E increases until it reaches 0.7 and then starts to decrease when we further increase λ_E . However, for

Table 5: Ranking performance of different variations of SetRank. Best results are marked bold. The superscript “*” means the model significantly outperforms SetRank^{-ts} (with p-value ≤ 0.05).

Dataset	Metric	SetRank ^{-ts}	SetRank ^{-t}	SetRank
S2-CS	NDCG@5	0.3847	0.4157*	0.4207*
	NDCG@10	0.4095	0.4423*	0.4431*
	NDCG@15	0.4256	0.4655*	0.4762*
	NDCG@20	0.4443	0.4813*	0.4950*
TREC-BIO	NDCG@5	0.3414	0.3705	0.3744
	NDCG@10	0.3257	0.3500	0.3522*
	NDCG@15	0.3140	0.3335	0.3363*
	NDCG@20	0.3058	0.3217	0.3246
S2-CS -ESQ	NDCG@5	0.4059	0.4800*	0.4983*
	NDCG@10	0.4311	0.5004*	0.5130*
	NDCG@15	0.4469	0.5266*	0.5450*
	NDCG@20	0.4683	0.5378*	0.5629*
TREC-BIO -ESQ	NDCG@5	0.3257	0.3594	0.3639*
	NDCG@10	0.3100	0.3380*	0.3406*
	NDCG@15	0.2994	0.3219*	0.3251*
	NDCG@20	0.2903	0.3100*	0.3132*

Figure 4: Sensitivity of λ_E in S2-CS and TREC-BIO datasets.

TREC-BIO dataset, the optimal value of λ_E is around 0.3, and if we increases λ_E over 0.6, the ranking performance will drop quickly.

5.4 Effectiveness of Model Selection

5.4.1 Experimental Setup.

In this experiment, we try to apply our unsupervised model selection algorithm to choose the best parameter settings of SetRank without using a validation set. We select entity token weight λ_E , title field weight δ_{title} , abstract field weight δ_{abs} , dirichlet smoothing factors for both fields μ_{title} & μ_{abs} from $\{0.2, 0.3, \dots, 0.8\}$, $\{5, 10, 15, 20\}$, $\{1, 3, 5, 10\}$, and $\{500, 1000, 1500, 2000\}$, respectively. This generates totally $7 \times 4 \times 4 \times 4 \times 4 = 1,792$ possible parameter settings and for each of them we can construct a ranking model. We first apply our unsupervised model selection algorithm (with either *KT* or *posKT* as the ranking distance) and obtain the most confident parameter settings returned by it. Then, we plug in these parameter settings into SetRank and denote it as AutoSetRank. For reference, we also calculate the average performance of all 1,792 ranking models.

Table 6: A case study comparing SetRank with SetRank^{-ts} on one entity-set query in S2-CS. Note: Atari is a video game platform.

Query	reinforcement learning for video game	
Method	SetRank ^{-ts}	SetRank
1	The effects of video game playing on attention, memory, and executive control	A video game description language for model-based or interactive learning
2	Can training in a real-time strategy video game attenuate cognitive decline in older adults?	Playing Atari with Deep Reinforcement Learning
3	A video game description language for model-based or interactive learning	Real-time neuroevolution in the NERO video game

Table 7: Effectiveness of ranking model selection. SetRank-VS: parameters are tuned using 5-fold cross validation. AutoSetRank-(KT/posKT): parameters are obtained based on our unsupervised model selection algorithm, which uses either KT or posKT as ranking distance. Mean ($\pm$ Std): the averaged performance of all ranking models with standard derivation shown.

Dataset	Method	δ_{title}	δ_{abs}	λ_E	μ_{title}	μ_{abs}	NDCG@5	NDCG@10	NDCG@15	NDCG@20
S2-CS	SetRank-VS	20	5	0.7	1000	1000	0.4207	0.4431	0.4762	0.4950
	AutoSetRank-KT	20	7	0.7	1500	2000	0.4174	0.4427	0.4730	0.4929
	AutoSetRank-posKT	20	5	0.7	1500	1500	0.4173	0.4436	0.4731	0.4923
	Mean ($\pm$ Std)	-	-	-	-	-	0.3898 ($\pm$ 0.0112)	0.4128 ($\pm$ 0.0106)	0.4411 ($\pm$ 0.0161)	0.4543 ($\pm$ 0.0163)
	SetRank-VS	20	5	0.2	1000	1000	0.3744	0.3522	0.3363	0.3246
TREC-BIO	AutoSetRank-KT	20	5	0.2	1500	1000	0.3692	0.3472	0.3305	0.3173
	AutoSetRank-posKT	20	7	0.2	1000	1000	0.3748	0.3564	0.3367	0.3253
	Mean ($\pm$ Std)	-	-	-	-	-	0.3479 ($\pm$ 0.0103)	0.3238 ($\pm$ 0.0079)	0.3199 ($\pm$ 0.0079)	0.3036 ($\pm$ 0.0093)

5.4.2 Experimental Result and Analysis.

Table 7 shows the results, including the SetRank’s performance when a labeled validation set is given. First, we notice that for S2-CS dataset, although the parameter settings tuned over validation set do perform better than the ones returned by our unsupervised model selection algorithm, the difference is not significant. For TREC-BIO dataset, it is surprising to find that AutoSetRank-posKT can slightly outperforms SetRank tuned on validation set. Furthermore, the performance of AutoSetRank function is higher than the average performances of all possible ranking models by 2 standard deviations, which demonstrates the effectiveness of our unsupervised model selection algorithm.

5.5 Use Case Study: Bio-Literature Search

In this section we demonstrate the effectiveness of SetRank in a biomedical use case. As preparation, we build a biomedical literature search engine based on over 27 million papers retrieved from PubMed. Entities in all papers are extracted and typed using PubTator. This search system is cold-started with our proposed SetRank model and we show how SetRank can help this search system to accommodate a given entity-set query and returns a high-quality rank list of papers relevant to the query. Comparison with PubMed, a widely used search engine for biomedical literature, will also be discussed.

A biomedical case. Consider the following case of a biomedical information need. Genomics studies often identify sets of genes as having important roles to play in the processes or conditions under investigation, and the investigators seek to understand better what biological insights such a list of genes might provide. Suppose such a study, having examined brain gene expression patterns in old mice, identifies ten genes as being of potential interest. The investigator forms a query with these 10 genes, submits it to a literature search engine, and examines the top ten returned papers to look for an association between this gene set and a disease. The query consists of symbols of the 10 genes: “APP, APOE, PSEN1, SORL1, PSEN2, ACE, CLU, BDNF, IL1B, MAPT”.

Relevance criterion. We choose the above ten genes for our illustration because these are actually top genes associated with Alzheimer’s disease according to DisGeNET [24], and it is unlikely

that there is another completely different (and unknown) commonality among them. Therefore, a retrieved paper is relevant if and only if it discusses at least one of the query genes in the context of Alzheimer’s disease. Furthermore, among all relevant papers, we prefer those covering more unique genes.

Result analysis. The top-5 papers returned by PubMed⁴ and our system are shown in Table 8. We see that the “Alzheimer’s disease” is explicitly mentioned in the title of all the five papers returned by our system, and the top two papers cover 6 unique genes among the total 10 genes. All five papers returned by SetRank are highly relevant, since they all focus on the association between a subset of our query genes and Alzheimer’s disease. In contrast, the top-5 papers retrieved by PubMed are dominated by two genes (i.e., APOE4 and BDNF) and contain none of the remaining eight. Only the 1st of the five papers is highly relevant. It focuses on the association between Alzheimer’s disease (mentioned explicitly in the title) and our query gene set. Three other papers (ranked 2nd to 4th) are marginally relevant, in the sense that Alzheimer’s disease is the context but not the focus of their studies. The paper ranked 5th is irrelevant. Therefore, users will prefer SetRank since it returns papers covering a large-portion of an entity-set query and helps them to find the association between this entity set with Alzheimer’s disease.

6 CONCLUSIONS AND FUTURE WORK

In this paper, we study the problem of searching scientific literature using entity-set queries. A distinctive characteristic of entity-set queries is that they reflect user’s interest in inter-entity relations. To capture such information need, we propose SetRank, an unsupervised ranking framework which explicitly models entity relations among the entity set. Second, we develop a novel unsupervised model selection algorithm based on weighted rank aggregation to select SetRank’s parameters without relying on a labeled validation set. Experimental results on two benchmark datasets corroborate the effectiveness of SetRank and the usefulness of our model selection algorithm. We further discuss the power of SetRank with a real-world use case of biomedical literature search.

As a future direction, we would like to explore how we can go beyond pairwise entity relations and integrate higher-order

⁴Querying PubMed with the exact same query returns 0 document. To get reasonable results, PubMed users have to insert an OR logic between every pairs of genes, and change the default “sorting by most recent” to “sorting by best match”.

Table 8: A real-world use case comparing SetRank with PubMed. The input query contains a set of 10 genes and reflects user’s information need of finding an association between this gene set and an unknown disease. Entity mentions in returned paper titles are highlighted in brown and the entity mentions of Alzheimer’s disease, which are used to judge paper relevance, are marked in red.

Query	APP APOE4 PSEN1 SORL1 PSEN2 ACE CLU BDNF IL1B MAPT	
Method	Rank	Paper Title
PubMed	1	Apathy and APOE4 are associated with reduced BDNF levels in Alzheimer’s disease
	2	ApoE4 and Aβ Oligomers Reduce BDNF Expression via HDAC Nuclear Translocation
	3	Cognitive deficits and disruption of neurogenesis in a mouse model of apolipoprotein E4 domain interaction
	4	APOE-epsilon4 and aging of medial temporal lobe gray matter in healthy adults older than 50 years
	5	Influence of BDNF Val66Met on the relationship between physical activity and brain volume
SetRank	1	Investigating the role of rare coding variability in Mendelian dementia genes (APP, PSEN1, PSEN2, GRN, MAPT, and PRNP) in late-onset Alzheimer’s disease
	2	Rare Genetic Variant in SORL1 May Increase Penetrance of Alzheimer’s Disease in a Family with Several Generations of APOE-4 Homozygosity
	3	APP, PSEN1, and PSEN2 mutations in early-onset Alzheimer disease: A genetic screening study of familial and sporadic cases
	4	Identification and description of three families with familial Alzheimer disease that segregate variants in the SORL1 gene
	5	The PSEN1, p.E318G variant increases the risk of Alzheimer’s disease in APOE-4 carriers

entity relations into the current SetRank framework. Besides, it would be interesting to explore whether SetRank can effectively model domain expert’s prior knowledge about the relative importance of entity relations. Furthermore, the incorporation of user interaction and and extension of current SetRank framework to weakly-supervised settings are also interesting research problems.

ACKNOWLEDGEMENTS

This research is sponsored in part by U.S. Army Research Lab. under Cooperative Agreement No. W911NF-09-2-0053 (NSCTA), DARPA under Agreement No. W911NF-17-C-0099, National Science Foundation IIS 16-18481, IIS 17-04532, and IIS-17-41317, DTRA HDTRA11810026, and grant 1U54GM114838 awarded by NIGMS through funds provided by the trans-NIH Big Data to Knowledge (BD2K) initiative (www.bd2k.nih.gov).

REFERENCES

- [1] Joran Beel and Bela Gipp. 2009. Google Scholar’s ranking algorithm: An Introductory Overview. In *ISSI*.
- [2] Avradeep Bhowmik and Joydeep Ghosh. 2017. LETOR Methods for Unsupervised Rank Aggregation. In *WWW*.
- [3] Pavel Brazdil and Christophe Giraud-Carrier. 2017. Metalearning and Algorithm Selection: progress, state of the art and introduction to the 2018 Special Issue. *Machine Learning* (2017).
- [4] David Carmel, Ming-Wei Chang, Evgeniy Gabrilovich, Bo-June Paul Hsu, and Kuansan Wang. 2014. ERD’14: entity recognition and disambiguation challenge. *SIGIR Forum* 48 (2014), 63–77.
- [5] Stéphane Clinchant and Éric Gaussier. 2010. Information-based models for ad hoc IR. In *SIGIR*.
- [6] Don Coppersmith, Lisa Fleischer, and Atri Rudra. 2006. Ordering by weighted number of wins gives a good ranking for weighted tournaments. In *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*. Society for Industrial and Applied Mathematics, 776–782.
- [7] Jeff Dalton, Laura Dietz, and James Allan. 2014. Entity query feature expansion using knowledge base links. In *SIGIR*.
- [8] Matthias Feurer, Aaron Klein, Katharina Eggenberger, Jost Tobias Springenberg, Manuel Blum, and Frank Hutter. 2015. Efficient and Robust Automated Machine Learning. In *NIPS*.
- [9] Edward A. Fox and Joseph A. Shaw. 1993. Combination of Multiple Searches. In *TREC*.
- [10] Dario Garigliotti and Krisztian Balog. 2017. On Type-Aware Entity Retrieval. In *ICTIR*.
- [11] John Guiver and Edward Snelson. 2009. Bayesian inference for Plackett-Luce ranking models. In *ICML*.
- [12] Jiafeng Guo, Gu Xu, Xueqi Cheng, and Hang Li. 2009. Named entity recognition in query. In *SIGIR*.
- [13] Faegheh Hasibi, Krisztian Balog, and Svein Erik Bratsberg. 2015. Entity linking in queries: Tasks and evaluation. In *Proceedings of the 2015 International Conference on The Theory of Information Retrieval*. ACM, 171–180.
- [14] William R. Hersh, Aaron Cohen, Jianji Yang, Ravi Teja Bhupatiraju, Phoebe Roberts, and Marti Hearst. 2005. TREC 2005 Genomics Track Overview. In *TREC*.
- [15] Sarvnaz Karimi, Justin Zobel, and Falk Scholer. 2012. Quantifying the impact of concept recognition on biomedical information retrieval. *Information Processing & Management* 48, 1 (2012), 94–106.
- [16] Maurice G Kendall. 1955. Rank correlation methods. (1955).
- [17] Alexandre Klementiev, Dan Roth, and Kevin Small. 2007. An Unsupervised Learning Algorithm for Rank Aggregation. In *ECML*.
- [18] Alexandre Klementiev, Dan Roth, and Kevin Small. 2008. A Framework for Unsupervised Rank Aggregation. In *SIGIR LR4IR Workshop*.
- [19] Xitong Liu and Hui Fang. 2015. Latent entity space: a novel retrieval approach for entity-bearing queries. *Information Retrieval Journal* 18 (2015), 473–503.
- [20] Zhiyong Lu. 2011. PubMed and beyond: a survey of web tools for searching biomedical literature. In *Database*.
- [21] Lucas Maystre and Matthias Grossglauser. 2015. Fast and Accurate Inference of Plackett-Luce Models. In *NIPS*.
- [22] Donald Metzler and W Bruce Croft. 2005. A Markov random field model for term dependencies. In *SIGIR*.
- [23] Paul Ogilvie and James P. Callan. 2003. Combining document representations for known-item search. In *SIGIR*.
- [24] Janet Piñero, Alex Bravo, Núria Queralt-Rosinach, Alba Gutiérrez-Sacristán, Jordi Deu-Pons, Emilio Centeno, Javier García-García, Ferran Sanz, and Laura I Furlong. 2016. DisGeNET: a comprehensive platform integrating information on human disease-associated genes and variants. *Nucleic acids research* (2016).
- [25] Hadas Raviv, Oren Kurland, and David Carmel. 2016. Document Retrieval Using Entity-Based Language Models. In *SIGIR*.
- [26] Xiang Ren, Jiaming Shen, Meng Qu, Xuan Wang, Zeqiu Wu, Qi Zhu, Meng Jiang, Fangbo Tao, Saurabh Sinha, David Liem, Peipei Ping, Richard M. Weinshilboum, and Jiawei Han. 2017. Life-iNet: A Structured Network-Based Knowledge Exploration and Analytics System for Life Sciences. In *ACL*.
- [27] Stephen E. Robertson and Hugo Zaragoza. 2009. The Probabilistic Relevance Framework: BM25 and Beyond. *Foundations and Trends in Information Retrieval* (2009).
- [28] Jiaming Shen, Zhenyu Song, Shitao Li, Zhaowei Tan, Yuning Mao, Luoyi Fu, Li Song, and Xinbing Wang. 2016. Modeling Topic-Level Academic Influence in Scientific Literatures. In *AAAI Workshop: Scholarly Big Data*.
- [29] Arnab Sinha, Zhihong Shen, Yang Song, Hao Ma, Darrin Eide, Bo-June Paul Hsu, and Kuansan Wang. 2015. An Overview of Microsoft Academic Service (MAS) and Applications. In *WWW*.
- [30] Jie Tang, Jing Zhang, Limin Yao, Juan-Zi Li, Li Zhang, and Zhong Su. 2008. ArnetMiner: extraction and mining of academic social networks. In *KDD*.
- [31] Christophe Van Gysel and Maarten de Rijke. 2018. Pytrec_eval: An Extremely Fast Python Interface to trec_eval. In *SIGIR*. ACM.
- [32] Chih-Hsuan Wei, Hung-Yu Kao, and Zhiyong Lu. 2013. PubTator: a web-based text mining tool for assisting biocuration. In *Nucleic Acids Research*.
- [33] Jian Wu, Kyle Williams, Hung-Hsuan Chen, Madian Khabsa, Cornelia Caragea, Alexander Ororbia, Douglas Jordan, and C. Lee Giles. 2014. CiteSeerX: AI in a Digital Library Search Engine. *AI Magazine* 36 (2014), 35–48.
- [34] Chenyan Xiong and James P. Callan. 2015. EsdRank: Connecting Query and Documents through External Semi-Structured Data. In *CIKM*.
- [35] Chenyan Xiong and James P. Callan. 2015. Query Expansion with Freebase. In *ICTIR*.
- [36] Chenyan Xiong, James P. Callan, and Tie-Yan Liu. 2016. Bag-of-Entities Representation for Ranking. In *ICTIR*.
- [37] Chenyan Xiong, James P. Callan, and Tie-Yan Liu. 2017. Word-Entity Duet Representations for Document Ranking. In *SIGIR*.
- [38] Chenyan Xiong, Russell Power, and James P. Callan. 2017. Explicit Semantic Ranking for Academic Search via Knowledge Graph Embedding. In *WWW*.
- [39] Yang Xu, Gareth J. F. Jones, and Bin Wang. 2009. Query dependent pseudo-relevance feedback based on wikipedia. In *SIGIR*.
- [40] ChengXiang Zhai and John D. Lafferty. 2001. A Study of Smoothing Methods for Language Models Applied to Ad Hoc Information Retrieval. *SIGIR Forum* 51 (2001), 268–276.
- [41] Zhibing Zhao, Peter Piech, and Lirong Xia. 2016. Learning Mixtures of Plackett-Luce Models. In *ICML*.