Information Sciences 384 (2017) 1–20

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

CI-Rank: Collective importance ranking for keyword search in databases

Xiaohui Yu[a,b,*], Ziqiang Yu[c,*], Yang Liu[a], Huxia Shi[b]

[a] *School of Computer Science and Technology, Shandong University, Jinan, Shandong 250101, China*
[b] *School of Information Technology, York University, Toronto, ON M3J 1P3, Canada*
[c] *School of Information Science and Engineering, University of Jinan, Jinan, Shandong 250022, China*

ARTICLE INFO

Article history:
Received 4 February 2015
Revised 16 October 2016
Accepted 13 December 2016
Available online 18 December 2016

Keywords:
Keyword search
Ranking
Random walk

ABSTRACT

Keyword search over databases, popularized by keyword search in WWW, allows ordinary users to access database information without the knowledge of structured query languages and database schemas. Most of the previous studies in this area use IR-style ranking, which fail to consider the importance of the query answers. In this paper, we propose CI-RANK, a new approach to keyword search in databases, which considers the importance of individual nodes in a query answer and the cohesiveness of the result structure in a balanced way. CI-RANK is built upon a carefully designed model called Random Walk with Message Passing that helps capture the relationships between different nodes in the query answer. This model lends itself well to adaptation to user preferences. We develop a branch and bound algorithm to support the efficient generation of top-k query answers. Indexing methods are also introduced to further speed up the run-time processing of queries. Extensive experiments conducted on two real data sets with a real user query log confirm the effectiveness and efficiency of CI-RANK.

© 2016 Elsevier Inc. All rights reserved.

1. Introduction

Recent years have seen a growing interest in supporting keyword search in databases [1–8]. Popularized by search engines on the Web, keyword search provides an intuitive, convenient, and effective way for users to interact with and explore the structured data stored in database. Different than the traditional ways of querying a database, it does not require the users to be familiar with the database schema and query languages.

A keyword query Q is specified by a set of keywords $Q = \{k_1, k_2, \ldots, k_{|Q|}\}$. Since those $|Q|$ keywords could find matches in different places in the database, the results are usually interconnected structures that contain all (or as many as possible) of the given keywords. For example, one popular form of such structures is joined tuple trees [7] for keyword search in relational database systems.

Conceptually, a database R can be viewed as a data graph $\mathbb{G}_R = (\mathbb{V}, \mathbb{E})$, where the vertex set $\mathbb{V}$ represents the data objects (e.g., tuples in the case of relational database systems), and $\mathbb{E}$ represents the connections between the objects (e.g., primary key-foreign key relationships). In general, keyword search could yield more than one result, as there could be multiple objects matching each keyword, and there could be different ways to assemble all matching tuples into a connected

---

[*] Corresponding author.
E-mail addresses: xhyu@yorku.ca (X. Yu), ise_yuzq@ujn.edu.cn (Z. Yu), yliu@sdu.edu.cn (Y. Liu), huxiashi@cse.yorku.ca (H. Shi).

http://dx.doi.org/10.1016/j.ins.2016.12.022
0020-0255/© 2016 Elsevier Inc. All rights reserved.

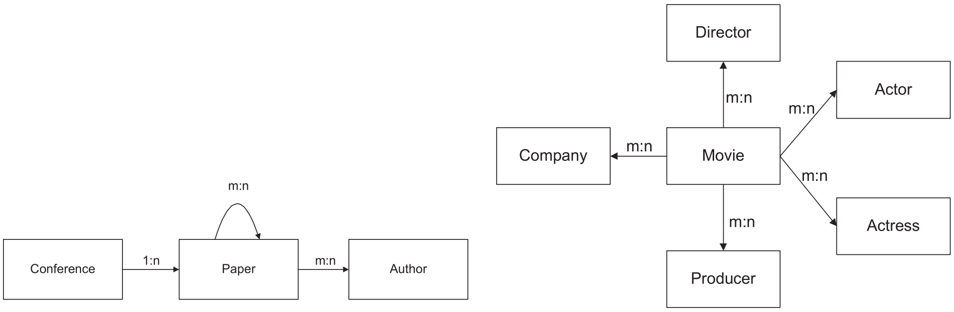

(a) DBLP (b) IMDB

Fig. 1. The schemas of the DBLP and IMDB data.

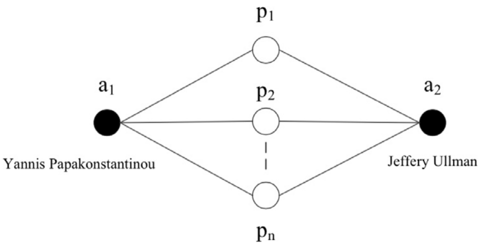

Fig. 2. Search results for the query "`Papakonstantinou Ullman`". p_1: Capability based mediation in TSIMMIS; p_2: The TSIMMIS Project: Integration of heterogeneous information sources.

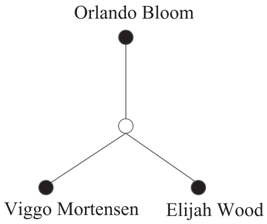

Fig. 3. The answer to the query "`Bloom Wood Mortensen`".

structure. Keyword search algorithms usually use a function *score*(·) to assign a score to a structure, and rank the set of results accordingly.

Various factors have been considered in designing the *score*(·) function. Early approaches simply use the size of the connected structure as the metric [4]. Many existing approaches utilize IR-style scoring functions [5–7]. The idea is to assign each object in the result an IR score based on its textual content, and then combine individual scores together using a score aggregation function *comb*(·) to obtain the final score. Other factors utilized in the literature include the distance between two keyword-matching objects in the connected structure [9], the coverage of keywords [7], etc.

Despite all recent efforts in improving the effectiveness of search through designing better scoring functions, one critical factor has been largely overlooked: the *importance* of the result structures. As an example, consider the following keyword query, "`Papakonstantinou Ullman`", on a bibliography database (DBLP) with the relational schema shown in Fig. 1(a).

The two keywords in the query match two tuples in the author table containing "Yannis Papakonstantinou" and "Jeffery Ullman" respectively. Each paper co-authored by them could connect the two tuples to form a complete result (Fig. 2). Hence there could potentially exist many results, which only differ in the paper tuple that connects the two authors together. When ranking the (potentially many) results generated, it is certainly desirable to assign higher scores to the results containing papers with more citations (thus considered more important). However, existing approaches do not take this importance factor into account. The widely used IR-style scoring functions only consider the textual information contained in the results, and thus cannot distinguish the more important results from the others. Since all results have the same size, incorporating the size factor in ranking the results does not help either (Fig. 3).

The importance of search results has been considered in ObjectRank [10,11]. However, the goal of that work is to rank individual objects (e.g., tuples) according to their relevance to the query. This in contrast to our task where we would like to measure the *collective importance* of the objects contained in a result. ObjectRank cannot be easily extended to handle this new task. BANKS [2] and Bi-directional Search [3] only consider the importance of the root node and the nodes containing the keywords in a result tree, and do not consider the importance of intermediate nodes which we will show in Section 2 has a significant impact on the quality of results. Measuring the importance of objects (such as Web documents) has also been well studied in the Web area, as exemplified by PageRank [12]. However, measuring the importance of keyword search results is a significantly different task and presents unique challenges. Whereas Web documents are single objects, keyword search results are usually connected structures. In addition, Web documents are entities that already exist and their importance values can be computed offline, while for keyword search in databases, the results are obtained dynamically based on the keywords and thus their importance values have to be computed on the fly to return the top-k results to users as soon as possible.

In this paper, we propose a novel method called CI-Rank (for Collective Importance Ranking) to rank keyword search results based on their collective importance. Note that the collective importance of a search result is not simply an aggregate of the importance of individual nodes, it is also affected by the structure of the result (i.e., how those nodes are connected together). CI-Rank takes into consideration not only the importance of the objects in the result, but also the structure of the result. To the best of our knowledge, this is the first work that addresses the issue of evaluating the overall importance of a search result. For ease of presentation, we focus on relational databases where the search results are joined tuple trees, but our approach is general enough to be applied to other types of structured data that can be modeled as graphs, such as XML data.

The underpinning of CI-Rank is a model called RWMP (for Random Walk with Message Passing). Our proposal of RWMP is inspired by the random walk model, which has been successfully employed in $\mathbb{G}_D$ [11] as well as in PageRank to compute the importance of individual nodes. The main innovation of RWMP is that it incorporates a carefully designed message passing procedure that helps model the degree of cohesiveness in the result tree in addition to the importance of the individual nodes.

CI-Rank defines a new scoring function based on RWMP. This new function calls for an algorithm that can produce the search results efficiently. To this end, we proposed a branch-and-bound algorithm that utilizes the properties of the scoring function to help prioritize the search and prune the search space. An index structure is also proposed to speedup the runtime computation through pre-computation and materialization of partial results.

The main contributions of this paper can be summarized as follows.

- We propose CI-Rank, a novel method for ranking keyword search results, which considers both the importance of nodes and the cohesiveness of the result tree structure. CI-Rank is enabled by our proposal of RWMP, a random walk model with message passing.
- We present a branch-and-bound tree search algorithm for computing the keyword search results, which allows for effectiveness pruning of the search space, while at the same time guarantees the optimality of the results.
- We design two index structures (star indexing and group indexing) to further improve the search efficiency, and utilize the characteristics of the database graph to reduce the sizes of two indexes effectively.
- We show how CI-Rank can adapt to user preferences using user feedback, and demonstrate the superiority of the ranking strategy employed by CI-Rank to a PageRank-style ranking function.
- We perform extensive experiments using both synthetic and real data and queries, including a manually labeled AOL query log, to demonstrate the superiority of CI-Rank over existing methods.

The rest of the paper is organized as follows. Section 2 defines the problem and analyzes problems with existing ranking methods. Section 3 presents CI-Rank and the random walk model with message passing. Section 4 describes the branch and bound algorithm for generating top-*k* query answers. The indexing techniques are presented in Section 5. Experimental results are presented in Section 6. Section 7 discusses related work. Section 8 concludes this paper and discusses possible directions for future work.

2. Preliminaries

In this section, we first formally define the keyword search problem, and then provide an analysis of the problems with existing methods for ranking search results.

2.1. Problem definition

A database is modeled as a weighted directed graph $\mathbb{G} = (\mathbb{V}, \mathbb{E})$. Each tuple t_i is represented by a node $v_i \in \mathbb{V}$. For any two nodes v_i, v_j, there is a directed edge $\langle v_i, v_j \rangle$ (or $v_i \rightarrow v_j$) and a backward edge $\langle v_j, v_i \rangle$ (or $v_j \rightarrow v_i$) if and only if there exists a foreign key on tuple t_i that refers to the primary key tuple v_j. The reason to model a foreign key to primary key relation as two directed edges rather than one undirected edge is to reflect the difference between these two directions. For instance, if paper p_b cites paper p_a, readers of p_b are more likely to also read p_a, whereas readers of p_a are less likely to also read p_b.

In the following discussion, we assume the AND semantics between the keywords (i.e., the results should contain all keywords).

Definition 1 (input query)**.** An input *query* Q is a set of keywords, $Q = \{k_1, k_2, \ldots k_{|Q|}\}$, where k_i is keyword.

Definition 2 (free and non-free nodes [4])**.** If a database tuple matches any keyword k_i in the query Q, it is called a *non-free tuple* w.r.t. k_i. The corresponding node in graph is called a *non-free node* w.r.t. k_i. If a tuple does not contain the query keyword k_i, it is a *free tuple* w.r.t. k_i and the corresponding node is a *free node* w.r.t. k_i. We use $E_n(k_i)$ and $E_f(k_i)$ to denote the non-free and free node sets of k_i. From their definitions, we have $E_n(k_i) \cup E_f(k_i) = \mathbb{V}$ and $E_n(k_i) \cap E_f(k_i) = \phi$. $E_n(Q)$ is the set of all non-free nodes of the query Q. All nodes that contain any keyword in Q belong to this set. Then, $E_n(Q) = E_n(k_1) \cup E_n(k_2) \cup \ldots \cup E_n(k_{|Q|})$. Similarly, $E_f(Q) = E_f(k_1) \cap E_f(k_2) \cap \ldots \cap E_f(k_{|Q|})$ denotes the free node set of the query Q. Each node in this set does not contain *any* query keyword in Q.

Definition 3 (query answer)**.** An answer to query Q is defined as any subtree T of $\mathbb{G}_D$ such that T is reduced w.r.t. Q. That is, there exist a set of nodes in T, $\mathbb{R} = \{v_1, \ldots, v_m\}$, where v_i contains keyword k_i ($1 \le i \le m$), such that the leaves of T must come from $\mathbb{R}$, $leaves(T) \subseteq \mathbb{R}$, and the root of T must be also from $\mathbb{R}$ if the root has only one child. We call such a query answer a *joined tuple tree* (JTT).

Since there could be multiple JTTs for a query Q, a core task of keyword search in databases is to perform effective ranking of the JTTs.

2.2. Problems with existing ranking methods

2.2.1. IR-Style Ranking functions

Most existing solutions employ IR-style ranking. For example, DISCOVER2 [5] uses the following scoring function based on TF-IDF:

$$score(T,Q) = \frac{\sum_{v\in T} score(v,Q)}{size(T)},$$

$$score(v,Q) = \sum_{k\in v\cap Q} \frac{1+\ln(1+\ln(tf_k(v)))}{(1-s)+s\cdot\frac{dl_v}{avdl_v}} \cdot \ln(idf_k),$$

where $idf_k = \frac{N_{Rel(v)}+1}{df_k(Rel(v))}$, $tf_k(v)$ is the number of occurrences of keyword k in node v, dl_v denotes the length of the text in v, and $avdl_t$ is the average length of the text in the relation which the tuple corresponding to v belongs to (i.e., $Rel(v)$), $N_{Rel(v)}$ is the number of tuples in $Rel(v)$, $df_k(Rel(v))$ is the number of tuples in $Rel(v)$ containing keyword k, and s is a constant.

As another example, SPARK [7], the state-of-the-art IR-style method, considers three factors, $score_a(T, Q)$, $score_b(T, Q)$, and $score_c(T, Q)$, where $score_a(T, Q)$ is the TF-IDF score, $score_b(T, Q)$ is called the completeness factor, and $score_c(T, Q)$ the size normalization factor. The final scoring function is defined as

$$score(T,Q) = score_a(T,Q) \cdot score_b(T,Q) \cdot score_c(T,Q)$$

The IR score of a JTT T is determined by

$$score_a(T,Q) = \sum_{k\in T\cap Q} \frac{1+\ln(1+\ln(tf_k(T)))}{(1-s)+s\cdot\frac{dl_T}{avdl_{CN^*(T)}}} \cdot \ln(idf_k),$$

where $tf_k(T) = \sum_{v\in T} tf_k(v)$, $idf_k = \frac{N_{CN^*(T)}+1}{df_k(CN^*(T))}$, and $CN^*(T)$ denotes the join of relations containing the keywords (see [7] for the precise definition), dl_T is the total length of all text attributes for tuples in T, and $avdl_{CN^*(T)}$ is the average length of all tuples in $CN^*(T)$.

The completeness factor measures the coverage of keywords by T to allow flexibility in forcing AND/OR semantics, and the size normalization factor normalizes the score of T based on its size to ensure fairness between small and large JTTs. The exact formulas for those two factors are omitted here due to the limited space.

The aforementioned two scoring functions only consider the textual information, and fail to consider the difference in the importance of the nodes. Continuing the example shown in Fig. 2, the nodes a_i and p_i represent an author and a paper respectively. For the query ”Papakonstantinou Ullman”, we find that two paper nodes p_1 and p_2 can both link the two matching author nodes together into two JTTs: $a_1 \bowtie p_1 \bowtie a_2$ and $a_1 \bowtie p_2 \bowtie a_2$ (where $\bowtie$ denotes a join). It can be observed from the data that paper p_1 is cited 7 times and paper p_2 38 times. Apparently, the JTT containing p_2 should be ranked higher than that containing p_1 as it is considered to be more important. However, if we use the DISCOVER2 scoring function, the two corresponding JTTs have exactly the same score, as the paper nodes p_1 and p_2 do not match the keywords and therefore do not affect the final scores. Using the SPARK scoring function does not help either: we can show that the score for the JTT ($a_1 \bowtie p_2 \bowtie a_2$) is actually *lower* than that for the JTT ($a_1 \bowtie p_1 \bowtie a_2$). The reason is that the only factor that differentiates the scores of those two JTTs is dl_T with all other factors in the scoring function being equal. Since p_1 has a shorter title than p_2, it has a smaller dl_T and thus a higher final score.

2.2.2. Graph-based ranking functions

The ranking functions of BANKS [2] and Bi-directional Search [3] combine two scores, the node and edge scores, to generate the overall tree score. The node score is the average weight of the root node and the leaf nodes. The edge score is $1/(1+sum(e))$, where $sum(e)$ is the total weight of all edges in the answer tree. However, this scoring method cannot address the difference in the importance of the free nodes. As an example, consider the query “`Bloom Wood Mortensen`” on the IMDB database with the schema shown in Fig. 1(b). The result JTT may have the Actor node “Orlando Bloom” as the root, “Elijah Wood” and “Viggo Mortensen” as the leaves, and any movie co-stared by the three actors as the intermediate free node. Assume that the weights of the edges are independent of the movie. Since only the weights of the root and leaf nodes are considered, choosing which movie to link them together does not make a difference in the final score. This is certainly undesirable, because more popular movies should be favored. This example illustrates that all nodes and their connections should be accounted for in scoring the query answer.

3. Measuring the collective importance of query answers

The analysis in the preceding section reveals that the collective importance of query answers cannot be properly measured by existing ranking functions and some alternatives we have considered. This motivates our proposal of CI-Rank, which not only considers the importance of individual nodes (both free and non-free), but also the connection structures between them. Inspired by PageRank, we propose a variant of the random walk model called Random Walk with Message Passing (RWMP) to provide the basis for the design of our scoring function. In the remainder of this section, we will first describe the classic random walk model, and then present RWMP. Finally, we discuss how to score query answers with RWMP.

3.1. The random walk model

The random walk model has been successfully applied in computing PageRank scores of Web pages [12,13]. The whole Web can be considered as a graph $(\mathbb{V}, \mathbb{E})$, where a Web page corresponds to a node in $\mathbb{V}$, and a hyperlink is represented as an edge in $\mathbb{E}$. A random surfer moves in the graph from node to node in two ways. In any step, the random surfer currently in node v_i can either fly to a random node in the graph with probability c, or walk to a neighbor node with probability $1 - c$. $c \in (0, 1)$ is called the “teleportation” constant.

Let p_i denote the probability of node v_i being visited by the random surfer. It measures the importance of that node and is considered the PageRank value of the corresponding Web page. The vector $\mathbf{p} = \{p_1, p_2, \dots p_i, \dots, p_{|\mathbb{V}|}\}$ can be obtained by solving the following equation.

$$\mathbf{p} = (1 - c) \cdot \mathbf{Mp} + c \cdot \mathbf{u} \tag{1}$$

The first part $(1 - c) \cdot \mathbf{Mp}$ corresponds to a random surfer walking from a node to its neighbor, where $\mathbf{M}$ is the adjacency matrix with size $|\mathbb{V}| \times |\mathbb{V}|$. If there exists an edge $\langle v_j v_i \rangle$, m_{ij} is $1/OutDegree(v_j)$; m_{ij} is zero otherwise. The second part $c \cdot \mathbf{u}$ corresponds to the surfer flying to a random node. $\mathbf{u}$ is a probability vector called the “teleportation vector”, where u_i in $\mathbf{u}$ is the probability of node v_i being the destination the surfer is flying to.

The random walk model can also be applied in the database context to compute the importance of individual nodes. Treating the database R as a directed graph $\mathbb{G}_R$, the surfer moves within $\mathbb{G}_R$ in a similar fashion as described above. The edge weights can be determined and tuned as needed, and the importance of a node is the probability that a random surfer appears in this node in a steady state. This value can be computed by iteration or Monte Carlo simulation of Eq. (1).

3.2. Extending the random walk model for keyword search in databases

The random walk model works well in the Web context, but it cannot be applied directly to keyword search in databases, because the answers are trees instead of single nodes. Intuitively, we want the search results, i.e., the JTTs, to be generally important and at the same time tightly connected. Therefore, the scoring function must reflect both the importance of individual nodes and the cohesiveness of the result JTTs.

A naive way to extend the random walk model to database search is to score a JTT using the average importance values of the non-free nodes in it. The problem with this approach is that the cohesiveness of the JTT is not accounted for. For example, suppose for a keyword query $\{k_1, k_2\}$, both v_1 and v_3 match k_1, and both v_2 and v_4 match k_2. Let p_i be the importance value of v_i. If $avg(p_1, p_2) > avg(p_3, p_4)$, then the JTT containing v_1 and v_2 would be ranked higher than the JTT containing v_3 and v_4. However, if v_1 and v_2 are connected by a long path or there is no meaningful connection between the two nodes at all, while v_3 and v_4 are closely connected, the user might prefer the JTT containing v_3 and v_4.

Another possible scoring function is to take the average importance value of all nodes, free and non-free, in a JTT. This function also has a major flaw, which we call the *free node domination* problem. If the score of a JTT is dominated by the free nodes, the irrelevant JTTs that contain of irrelevant non-free nodes could be ranked higher than the relevant JTTs. For example, the answer to the query “`wilson cruz`” on the IMDB database with the schema shown in Fig. 1(b) should be a single node tree T_1 as shown in Fig. 4. However, another tree T_2, which is also shown in Fig. 4 and is obviously irrelevant, also matches this query. The non-free nodes are the movie “Charlie Wilson’s War (2007)” and the actress “Penelope Cruz”, respectively. They are connected by two free nodes, “Tom Hanks” and “America: A Tribute to Heroes (2001)”. Because the importance value of the free node “Tom Hanks” in T_2 is much higher that of the single node in T_1, the average importance value of T_2 is higher than T_1, and is therefore ranked higher. This illustrates the problem caused by domination of free nodes.

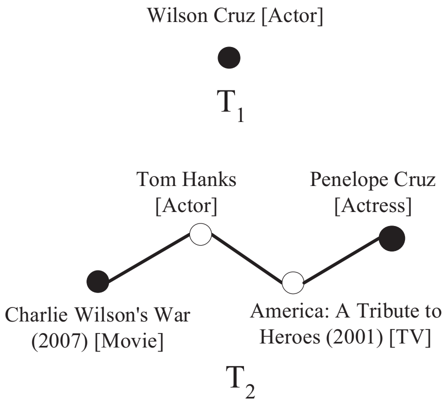

+0.2in+0.2in

Fig. 4. The free node domination problem.

Yet another alternative ranking function is to use *the average importance score / result size*, which takes not only the average importance of nodes, but the size of the result into consideration. The problem is that it cannot address the structural difference between the answer trees. For example, suppose we have two JTTs, T_1 and T_2, for the same query. T_1 has a star structure with one free node in the center and four non-free nodes connecting directly to the free node. T_2 also has one free node and four non-free nodes. They have exactly the same importance scores as those in T_1, but they form a chain-like structure with the free one in the center. Since T_1 and T_2 have the same size, and the importance values of the nodes are the same, T_1 and T_2 will have the same ranking scores despite their significant structural difference and hence potentially very different semantics.

3.3. Random walk with message passing

We propose a random walk model with message passing (RWMP) that is designed to match the characteristics of keyword search in databases. It considers both the importance of the single nodes and the structural cohesiveness of the answer trees in a balanced way, avoiding the pitfalls of the ranking functions discussed in the preceding subsection.

3.3.1. Overview of the model

The RWMP model consists of the following procedures.

1. Message generation: In RWMP, a non-free node generates messages indicating the existence of itself. The messages generated from different nodes have different types. We call such a message generated at v_i a v_i message. The number of messages generated at v_i is proportional the importance value of v_i. Let p_i be the importance value of the node v_i, and $|v_i|$ the number of words in node v_i. $|v_i \cap Q|$ denotes how many words in the node v_i match the query Q. The number of the message notes generated at node v_i is $r_{ii} = (t \cdot p_i \cdot |v_i \cap Q|)/|v_i|$, where t is the total number of the random surfers in the graph.

2. Message passing: One of the design objectives of RWMP is to measure the strength of pairwise connection of the non-free nodes in an answer tree. It is very difficult to measure the connection inside an answer tree using the original random walk model. We therefore add to it a component called message passing.

Intuitively, the message-passing model can be thought of as follows. A source node (which contains a keyword in the query) sends out a signal. The strength of the signal is weakened when it passes through an intermediate node. By measuring the strength of the received signal in the destination node, we know how closely these two nodes are connected.

To be more specific, the random surfers move to the neighboring nodes, carrying the messages. Since we are interested in evaluating the cohesiveness of the result tree, we assume that messages are only passed between nodes inside the tree and the surfers leaving the tree do not carry any messages. The number of surfers moving along different outgoing edges are proportional to the weights of those edges (which can be determined by domain experts or in ways similar to what is done in BANKS [2]).

Let the number of v_i messages received at node v_j be r_{ij}, and the number of leaving messages be f_{ij}. Suppose v_k is a neighbor of v_j and v_k is not in the path that connects v_i and v_j in the tree T. Then, at node v_k, the number of the v_i messages received, r_{ik}, is $f_{ij} \cdot w_{jk} / \sum_{v_n \in N(v_j) \cap V(T)} w_{jn}$, where w_{jk} is the weight of edge e_{jk}, $N(v_j)$ is the set of neighboring nodes of v_j in $\mathbb{G}_R$, $\sum_{v_n \in N(v_j) \cap V(T)} w_{jn}$ is the sum of the weights of all edges from v_j to its neighbors in the tree T. If the neighbor v_k is on the path that connects v_i and v_j in the tree T, v_k must have already got the v_i messages before v_j does. Since v_j sends v_i messages along all of its outgoing edges in the tree T, some v_i messages are sent back to v_k. These messages are discarded.

3. Message dampening: The number of messages is dampened in intermediate nodes, as surfers may drop with a certain probability some of the messages they are carrying. We call $d_j = f_{ij}/r_{ij} < 1 \quad (i \neq j)$ the *survival rate* at node v_j, representing the proportion of messages surviving the dampening process. The value of d_j is positively correlated with the importance of node v_j, and is the same across different message types. Apparently, the higher the survival rate is, the more messages are kept and passed on. We now describe in more detail the method for dampening messages and the choice of the survival rate.

3.3.2. Methods for dampening messages

Message dampening is an important procedure in RWMP as the dampening strength reflects the relative importance of intermediate nodes. Since we prefer the JTT to be connected by important nodes, the relationship between the survival rate d_i and the node importance p_i should be monotonically increasing. A straightforward design is $d_i \propto p_i$. However, its dampening strength is too heavy. In our experience, a high importance value can be thousands of times greater than a low one. As such, the range of the survival rates generated according to $d_i \propto p_i$ is too large and inflexible.

Our preliminary experimental results reveal that our model is highly effective when the survival rate of an intermediate node is proportional to the logarithm of its importance value. This design can be interpreted by the following in-node message exchange process. Suppose that in a node v_i, there are $t \cdot p_i$ random surfers. Each random surfer carrying messages talks to all other surfers inside this node in several steps:

(1) The surfer finds a group of g random surfers (regardless of the types of messages they carry) and talks to them. With a certain probability, she gives the message to one of the surfers in this group. We call this message-carrying surfer a speaker, and the group of g random surfers listeners.
(2) The listeners in the previous step become speakers. Each of them finds g listeners and talks to them with possible message exchange. A random surfer can be a listener only once in any step.
(3) Repeat this process, until all surfers in this node have communicated directly or indirectly with the given surfer.

In this manner, the surfer can communicate with all random surfers in this node in $log_g(t \cdot p_i)$ steps.

Let the probability that a surfer is willing to keep the messages by herself be α. With probability $1-\alpha$, she gives the messages to one of her listeners. In the last talk step, if the surfer who is now carrying the messages does not want to keep the messages, she just discards them. Hence, after $log_g(t \cdot p_i)$ talk steps, the probability that a message is not discarded is $\alpha \cdot \sum_{n=0}^{log_g(t \cdot p_i)} (1-\alpha)^n = 1-(1-\alpha)^{1+log_g(t \cdot p_i)}$. This is the dampening function at this node.

The dampening function has three parameters, the probability of keeping the messages α, the group size g, and the total number of random surfers t. Recall that the node importance p_i is obtained from Eq. (1). To simplify the dampening function, we assume that the node with the lowest importance value has only one random surfer. Let $p_{\min}$ denote the importance of this node. Then, $t = 1/p_{\min}$, and the survival rate can be written as follows:

$$d_i = 1-(1-\alpha)^{1+log_g(p_i/p_{min})} \tag{2}$$

3.3.3. Scoring of JTTs

In RWMP, a non-free node v_i receives messages generated from other non-free nodes. The number of the types of the incoming messages is $|E_n(Q) \cap V(T)| - 1$. We choose to use the number of messages of the least populous type in v_i as its score, i.e.,

$$score(v_i) = \min_{v_j \in E_n(Q) \cap V(T),\, j \neq i} (f_{ji}) \tag{3}$$

Intuitively, in a destination node (non-free node), if we take one message of each type (corresponding to each source node) and combine them, that can be considered complete knowledge of all sources. The number of messages of the least populous type determines the number of such message combinations, which measures how strongly this node is connected to the other parts of the answer tree. The scoring function in Eq. (3) reflects the number of such combinations.

The scoring function of CI-Rank for the whole tree T is calculated as the average of the scores of all non-free nodes.

$$score(T) = \frac{\sum\limits_{v_i \in E_n(Q) \cap V(T)} (score(v_i))}{|E_n(Q) \cap V(T)|} \tag{4}$$

3.3.4. An example of RWMP

We use an example in Fig. 5 to illustrate the process of RWMP model. This example shows how the “Orlando Bloom” messages are sent to other nodes in the tree. In Fig. 5(a), the messages are generated and carried out of node “Orlando Bloom”. Suppose the total amount is 1000 in this step.

Fig. 5 (b) shows the message dampening in the middle node. We choose the movie “The Lord of the Rings: The Fellowship of the Ring (2001)” as the free node in the middle. When the “Orlando Bloom” messages arrive at this node, some of them are discarded. The number of the messages is reduced from 1000 to 411. If this node has lower importance value, more messages are discarded.

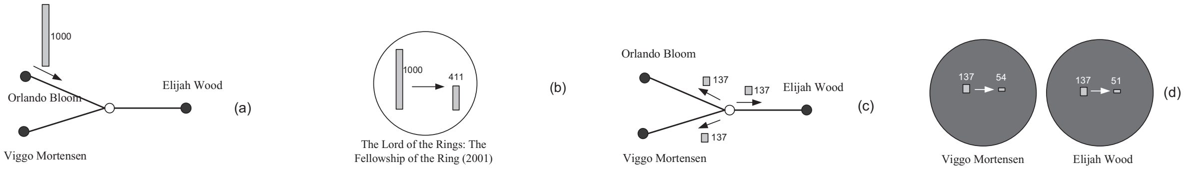

(a) 1000 messages leaving the source node corresponding to actor “Orlando Bloom”

(b) Message dampening in an intermediate node corresponding to the movie “Lord of the Rings”. Only 411 messages remain.

(c) Messages leaving the intermediate node, with 137 (=411/3) messages being sent along each edge

(d) Message dampening in the two destination nodes corresponding to the actors “Viggo Mortensen” and “Elijah Wood” respectively

Fig. 5. An example of RWMP.

In next step, the “Orlando Bloom” messages leave the middle movie node. This step is demonstrated in Fig. 5(c). Because the weights of the three out edges of the middle node are same, the out messages are equally distributed on them. Those messages that move back towards the node “Orlando Bloom” are ignored because they are messages about the node itself. The messages moving along the other two edges will enter the nodes “Viggo Mortensen” and “Elijah Wood”.

Inside these two nodes, these messages are dampened first as shown in Fig. 5(d). Then, they are recorded. For example, the node “Viggo Mortensen” will record that the received amount of the “Orlando Bloom” messages is 54.

Since all non-free nodes send out messages in a similar manner. Therefore, the node “Viggo Mortensen” also receives the “Elijah Wood” message. The number of these received messages are used to calculate the score of a single node and the whole JTT according to the formulas presented in previous subsection.

3.3.5. Benefits of CI-Rank

The benefits of the scoring model in CI-Rank can be summarized as follows.

3.4. Adapting to user preferences

The design of the proposed RWMP model allows it to be adapted to the user’s preferences through user feedback. Suppose there is a query log containing the past queries and the correct query answers chosen by the users. Biasing the scoring function in CI-Rank is done through the teleportation vector $\mathbf{u}$ in Eq. (1). We set u_i to be proportional to the frequency that its corresponding node v_i appears in the correct query answers. In this way, not only do the nodes that are contained in the correct query answers receive preference, the related nodes also get higher importance values.

The strength of this bias is adjustable by assigning different values to the constant c, the weight of the teleportation vector, in Eq. (1). In CI-Rank, initially the importance values of the database graph nodes can be calculated offline. When sufficient user feedback is accumulated in the form of query logs, the importance values of the nodes can be updated, again in an offline fashion. This is because the important values are query-independent, leading to minimal overhead at query processing time.

3.5. Semantics

We now discuss how to extend the default AND semantics to any combination of AND/OR semantics, with which an advanced user can enter any combination of input keywords connected by AND and OR operators. Our approach is to first translate such an input query to the conjunctive normal form (CNF), and then process the transformed queries, where a query consists of multiple clauses each corresponding to a disjunction of keywords in our problem. To this end, we define the transformational rules as follows:

- For a query with input keywords connected by only AND operators, each keyword will be deemed as a clause when translating this query into the CNF.
- If a query contains both AND and OR operators, the sequence of keywords linked by one or more OR operators is viewed as a clause, while each remaining keyword linked by AND operators corresponds to a clause.

Based on the transformation rules, we can translate any user input query into CNF. For instance, for a given query “*frequent* AND *pattern* AND *mining* OR *discover*”, on the basis of the above rules, it will be translated into CNF as $frequent \wedge pattern \wedge (mining \vee discover)$.

Now we give the formal definition of the input query based on CNF. An input *query* Q is a conjunction of clauses. A clause is called a *group* G_i, and $Q = G_1 \wedge G_2 \wedge \ldots \wedge G_n$. A group G_i is a disjunction of keywords. Let $|G_i|$ denote the length of the group G_i, which is the number of keywords it contains. The number of keywords in query Q is its length $|Q| = \sum_{i=1}^{n} |G_i|$.

Our model accepts two modes of user input. In the default mode, the users can simply input a sequence of keywords, and this sequence of keywords is interpreted using the default AND semantics. Or, for the advanced users, they can input any combination of keywords connected by AND and OR operators. Queries of this type will be translated to its conjunctive normal form (CNF) based on the transformational rules, and then handled as follows.

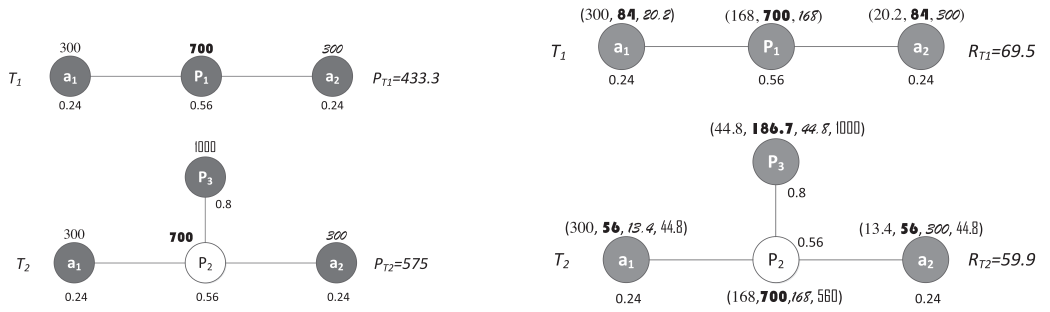

(a) T_1 and T_2 are two JTTs; P_{T_1} and P_{T_2} are the weights of T_1 and T_2 based on PageRankSearch

(b) The weight of each node with RWMP; R_{T_1} and R_{T_2} are the weights of T_1 and T_2 based on RWMP

Fig. 6. Comparison of CI-Rank and PageRankSearch.

For an input query Q in CNF, we require that any result of the query Q must contain at least one keyword from each group. In this case, a node matching any keyword in a group is considered a non-free node of this group, and a result refers to an optimal Steiner tree that covers all groups of Q. Hence, compared with handling the default AND semantics, the search algorithm (to be discussed in the next section) remains the same to process the input queries based on CNF, just by considering each group in the input as a single keyword in the default mode. Finally, we try to find top-k Steiner trees that cover all groups and have the largest scores.

3.6. Comparison with a PageRank-based ranking function

PageRank [14] and its variants have been widely used in web search and gained great success. So a natural alternative to consider is to apply PageRank to our setting. We thus design a ranking function called PageRankSearch based on the PageRank principle, and compare it with the ranking function employed by CI-Rank. PageRankSearch views the data graph representing the database as a graph of webpages (in the original PageRank context), where each node corresponds to a page and the edge represents a reference from one page to another, and employs a simplified algorithm [15] to compute the PageRank value of each node in an offline fashion. PageRankSearch defines the score of a *JTT* as the sum of the PageRank values of all the nodes in this tree divided by this tree size, i.e., $W(T) = \sum_{i=1}^{n} p(v_i)/Size(T)$, where $p(v_i)$ denotes the PageRank value of node v_i in a joined tuple tree *T*, and *Size*(*T*) is the number of nodes in *T*. The ranking of a given set of JTTs can be done efficiently at runtime as the PageRank values of the nodes can be computed offline.

However, the gain of PageRankSearch in efficiency comes at the cost of effectiveness. PageRankSearch fails to consider the structural cohesiveness of an answer tree, as it just summarizes the PageRank values of individual nodes. The nodes with large PageRank values will dominate the score of a JTT which leads to inferior performance. For instance, an answer tree that the user does not want could be ranked at the top when it contains a node with overly large PageRank value, regardless of whether this node is a free node or a non-free node. RWMP, on the other hand, avoids this problem as it makes the generated messages of every node pass through other nodes in a *JTT* with a survival rate, which reduces the impact of nodes with high importance values.

To be more concrete, let us consider the example in Fig. 6. In Fig. 6(a), T_1 and T_2 are two *JTTs*, where a_i (i=1, 2) and p_j (j=1, 2, 3) represent an author and a paper respectively. The green and white points represent non-free and free nodes respectively, and every node has a PageRank value and a survival rate that is introduced by RWMP. Take the node a_1 in T_1 as an instance. Its PageRank value and survival rate are 300 and 0.24 respectively. For comparison, we set the number of generated messages of each node in RWMP to be equal to its PageRank value. Fig. 6(b) shows the weights of all the nodes calculated using the RWMP model. In this example, T_1 represents the answer where a_1 and a_2 co-author p_1, while T_2 represents the answer where p_3 is cited by p_2 with co-authors a_1 and a_2. Obviously, T_1 reflects a more cohesive relationship between authors a_1 and a_2 and it is more likely to be a better answer. However, according to PageRankSearch, we have $P_{T_1} = 433.3$ and $P_{T_2} = 575$, which means that the importance of T_2 is higher than that of T_1. In contrast, we utilize RWMP to compute the weight of every node and the numbers with the same font represent the number of generated messages of a node being distributed in other nodes in a *JTT*. The results in Fig. 6(b) show that T_1 is more important than T_2 in this scenario ($R_{T_1} > R_{T_2}$), which coincides with our intuition.

4. Search algorithms

We now consider the problem of efficiently generating the top-*k* answers to the keyword queries. In many previous studies, the cost of an answer tree is the summation of its edge weights. In this case, the problem of generating the top-*k* answers becomes the minimum cost group Steiner tree problem, which has been shown to be NP-Complete. [16] lists several heuristic algorithms for this problem.

The scoring function in CI-Rank is more complex than the summation of the tree edge weights. Therefore, existing algorithms cannot be applied directly. In what follows, we will first describe a naive brute-forth algorithm, and then present a branch and bound search algorithm that utilizes the characteristics of the scoring function to get the top-k answers in a more efficient way. Similar to previous studies [2,4,5], we put a limit D on the diameter of answer trees.

4.1. The naive search algorithm

In the naive search algorithm, we first perform breadth first search from all non-free nodes. Then, the information gathered during the search is combined to generate the answer trees. The maximal search distance is $\lceil D/2 \rceil$ from each non-free node. When a node is visited in the search, the source node and the distance between the current node to the source node is recorded in this node. Because there may exist several different paths from the source node to this node being visited, the node visited right before this node is also recorded. After the search is finished for all non-free nodes, we check the visited nodes. If a node is reachable from one or more non-free nodes, and all keywords in the query are covered by these reachable non-free nodes, this node is selected as the root of the answer tree. From this root node, we connect the non-free nodes together to generate a full answer. Since there could be multiple paths from a root node to each non-free node, all combinations should be considered to generate different answers.

The main problem with this naive search algorithm is that it has to thoroughly expand all non-free nodes to find the best answer, a very time consuming task. To improve the efficiency, we design a branch and bound search algorithm, which prefers to expand the most promising non-free nodes and can stop early without exhaustively expanding all non-free nodes.

4.2. The branch and bound search algorithm

The branch and bound search algorithm works on candidate trees. A candidate tree C is a tree that covers at least one keyword. The initial candidate trees are the single non-free nodes constrained by the query keywords. We use *tree grow* and *tree merge* procedures [17] to expand small candidate trees to larger ones. The complete tree resulting from the expansion process that matches all input keywords, is considered an answer. At any time, only the top-k answers are kept.

The algorithm orders the candidate trees according to their upper bound values and expands the candidate tree with the highest upper bound value. Let us denote a candidate tree C with a root node v_i by $C(v_i)$. The upper bound of this candidate tree $C(v_i)$, $ub(C(v_i))$, is the maximal score that an answer tree can have if it is obtained from expanding this candidate tree $C(v_i)$. All candidate trees are kept in a priority queue P in the decreasing order of their upper bound values. The algorithm always retrieves the tree at the head of the queue P, the most promising candidate tree, and expands it. When the minimal score in the current top-k answers is higher than the upper bound value of the tree at the top of the queue, the algorithm can stop, because any answer trees generated in future steps are guaranteed to be worse than the top-k answers already obtained.

The outline of the branch and bound algorithm is shown in Algorithm 1.

Tree growing and tree merging: Let $v_j \in N(v_i)$ be a neighbor of v_i. If v_j is not contained in tree $C(v_i)$ (i.e., $v_j \notin C(v_i)$), we can create a new tree whose root is v_j with a single child tree $C(v_i)$. This process is called tree growing. $C(v_i)$ can merge with another candidate tree $C^{'}(v_i)$ if they have the same root, and the result covers more keywords than either $C(v_i)$ or $C^{'}(v_i)$. Invalid merge results with cycles are avoided through a sanity check.

Upper bound calculation: The tree growing and tree merging methods ensure the following important property of the candidate tree: when a candidate tree $C(v_i)$ is used as part of a larger tree $C(v_j)$, only the root of this small candidate tree, v_i, is connected to the other part of $C(v_j)$. $ub(C(v_i))$ is calculated according to this property from two different estimates, the *complete estimate* and the *potential estimate*, denoted by $ce(C(v_i))$ and $pe(C(v_i))$ respectively.

Complete estimate: The complete estimate $ce(C(v_i))$ is the highest possible score that an incomplete candidate tree $C(v_i)$ can achieve when it is expanded to a complete answer. An incomplete candidate tree $C(v_i)$ only covers a portion of the query keywords. To expand it to a complete tree, we need to supply non-free nodes that cover the missing keywords, and connect them by some paths to the root v_i of this candidate tree. If $C(v_i)$ is complete, the complete estimate is its own score $score(C(v_i))$. The difficulty of obtaining the complete estimate $ce(C(v_i))$ for an incomplete candidate tree $C(v_i)$ is that we do not know which non-free nodes can be used to complete the tree and how to choose the paths that connect them to the candidate tree. The fact that some parts of these paths can be shared makes the above decision even more difficult. To obtain a rough estimate, we assume that the best supplementary non-free nodes are connected directly to the root v_i of the candidate tree $C(v_i)$. This assumption leads to a complete tree that may not exist in reality. It is thus called a hypothetical tree.

The complete estimate given by the above method is usually much larger than the true value, and can be further refined by checking the direct neighbors of the root v_i. If a missing keyword can be covered by a node v_j that directly connects to v_i, we consider two scenarios. One is that v_j is contained in the hypothetical complete tree; the other is that this keyword is covered by a node not directly connected. In the latter case, we assume that the best supplement node is connected to the root v_i through its best neighbor. The greater estimate from these two scenarios is chosen as the complete estimate. This method provides a more accurate estimate, and is used in our experiments.

Potential estimate: Although a complete answer tree covers all input query keywords, it is still possible to add more non-free nodes to this tree to generate other complete trees with higher scores. The potential estimate, $pe(C(v_i))$, is the

Algorithm 1 Branch and bound search.

```
1:  Initialize the priority queue $P := \phi$
2:  Initialize the list of top-$k$ results $L_{top-k} := \phi$
3:  for each $v_i \in E_n(Q)$ do
4:      create a single node candidate tree $C(v_i)$
5:      enqueue $C(v_i)$ into queue $P$
6:  end for
7:  while $P \neq \phi$ do
8:      dequeue $P$ to $C(v_i)$
9:      if $|L_{top-k}| = k$ and $ub(C(v_i)) < minscore(L_{top-k})$ then
10:         return $L_{top-k}$
11:     end if
12:     $S_{grow}(C(v_i)) \leftarrow$ tree grow of $C(v_i)$
13:     for each $C(v_j) \in S_{grow}(C(v_i))$ do
14:         $L_{top-k} := L_{top-k} \cup \{C(v_j)\}$ if $C(v_j)$ is complete
15:         enqueue $C(v_j)$ into queue $P$
16:         $S_{merge}(C(v_j)) \leftarrow$ tree merge of $C(v_j)$
17:         for each $C(v_h) \in S_{merge}(C(v_j))$ do
18:             $L_{top-k} := L_{top-k} \cup \{C(v_h)\}$ if $C(v_h)$ is complete
19:             enqueue $C(v_h)$ into queue $P$
20:         end for
21:     end for
22:     if $|L_{top-k}| > k$ then
23:         truncate $L_{top-k}$ to size $k$
24:     end if
25: end while
```

maximal contribution from the additional nodes appended to a complete tree. The tree growing and tree merging method ensures that the additional nodes can only connect to a candidate tree $C(v_i)$ through its root v_i. Similar to the method to get the complete estimate $ce(C(v_i))$, we can consider the non-free nodes in the set of nodes consisting of v_i and all direct neighbors of v_i. All other non-free nodes are assumed to be connected to v_i through the best neighbor of v_i. When these additional nodes are added to the complete candidate tree $C(v_i)$, each of them will have a node score, which is defined in Eq. (3). We take the maximum among these node scores as the potential estimate $pe(C(v_i))$.

Combination of complete and potential estimates: The complete and potential estimates can be combined to give the upper bound for the candidate tree, $ub(C(v_i))$. The upper bound of a candidate tree $T(v_i)$ is the greater of the complete and potential estimates. This is stated in the following lemma.

Lemma 1 (Correctness of the upper bound)**.** *For an arbitrary full answer tree T grown from the candidate tree C,* $score(T) \leq ub(C) = \max(ce(C), pe(C))$.

Proof: See Appendix A.

Lemma 2 (Monotonicity of the upper bound)**.** *Suppose we have a candidate tree* $C(r_o)$. *This tree grows from its root* r_o *to a neighbor node* r_g, *and generates a new candidate tree* $C(r_g)$. *We have* $ub(C(r_g)) \leq ub(C(r_0))$.

Proof: See Appendix B.

Theorem 1 (Optimality of the search algorithm)**.** *The top-k results of the branch and bound search algorithm are guaranteed to be optimal. For any answer tree* $T \notin L_{top-k}$ *where* L_{top-k} *is the top-k results,* $score(T) \leq minscore(L_{top-k})$, *where* $minscore(L_{top-k})$ *is the minimal score of the top-k results.*

Proof. Prove by contradiction. If the top-k results L_{top-k} generated by the branch and bound search algorithm are not optimal, there should exist an answer tree T such that $score(T) > minscore(L_{top-k})$. Because this tree T is not generated by the branch and bound search algorithm, at least one of the candidate tree, denoted by C, that can be grown or merged to T is pruned. The pruning rule in Lines 24 and 25 of Algorithm 1 shows that $ub(C) \leq minscore(L_{top-k})$.

On the other hand, Lemma 1 states that $ub(C) \geq score(T)$. Considering the assumption of tree T, $score(T) > minscore(L_{top-k})$, we can get a result that $ub(C) > minscore(L_{top-k})$.

It is a contradiction that $ub(C) \leq minscore(L_{top-k})$ and $ub(C) > minscore(L_{top-k})$. Therefore, the initial assumption that the top-k results L_{top-k} are not optimal must be false. □

4.3. Computing scores of JTT incrementally

In the branch and bound algorithm, a candidate tree C will constantly evolve into a new tree when it absorbs a new node or is combined with another candidate tree. Once a new tree C^n is constructed, it is a challenging problem to computing the score of C^n efficiently. A naive method for calculating the score of C^n is to directly employ the RWMP model to let the message of each node in C^n go through all other nodes in this tree, and we can then get the number of messages of the least populous type in each node to calculate the final score of C^n. This method is straightforward to implement, but the main problem is that the message of each node needs to pass through the whole tree repeatedly when the tree is extended, which may result in prohibitive computational cost.

To improve the efficiency, we design ISF, an incremental approach inspired by the work by Rodger et al. [18] for computing scores of new trees based on their existing scores. In the branch and bound algorithm, the candidate tree C can be extended in two ways: adding a new node or being combined with another candidate tree. Hence, when computing the score of a new tree we need to take into account these two conditions.

Adding a new node: Suppose the tree C will become a new tree C^n by adding a new node v_x, and v_n is the adjacent node of v_x in C^n. If v_x is a free node, it will not generate messages and therefore will not affect the score of C^n. Thus, we just need to consider the case of v_x being a non-free node. In this case, ISF has to calculate the score of v_x and recompute scores of the nodes in C. In the ISF approach, we let each node record the number of all types of messages it has received. Since v_n is linked to v_x, each type of messages in C will reach v_x through v_n. Because v_n maintains the number of all types of messages it has received, we can immediately compute the number of each type of messages that v_x has received. Next, we let the messages generated by v_x pass through each node in C, and then update the number of messages of the least populous type of each node. Finally, we can compute the score of C^n by combining the scores of each node based on Eq. (4). As a special case, if the tree C^n covers all keywords, and the number of messages that are carried from v_x to v_n is larger than the number of messages generated by v_n, then the scores of the nodes in C will not change because the dampening rate of each node is fixed for different messages that can be inferred by Eq. (2). In this case, $Score(C^n) = (Score(C) + Score(v_x))/(En(Q) \cap V(C^n))$.

Being combined with another candidate tree: In the second case that C is merged with another candidate tree C^a through their common root, we have to consider two scenarios. If the new tree C^n constructed by merging C and C^a covers all keywords, then it is not necessary to exchange messages between C and C^a to calculate the score of C^n. Let v_r denote the common root, and v_m and v_n be the adjacent nodes of v_r in C and C^a respectively. Next, we only need to send the messages of the least populous type in v_m to C^a through the root v_r because other types of messages in v_m will not reduce the score of C^a. Similarly, only the messages of the least populous type in v_n will walk through v_r and the nodes in C. Since v_r and the nodes in C and C^a only receive one or two types of messages, they can update their scores instantly. Hence, C and C^a can get their new scores quickly. The score of C^n can be obtained by combining the new scores of v_r, C, and C^a based on Eq. (4).

If the new tree C^n does not cover all keywords, we then need to exchange all types of messages between C and C^a. That is, all types of messages in v_m and v_n will walk through each node in C^a and C according to their common root v_r with the RWMP model. In this case, each node in C^n needs to update their scores and then the score of C^n can be calculated with Eq. (4).

Overall, the ISF approach provides an incremental paradigm for computing the scores of trees that are continuously extended with the branch and bound algorithm. In the best case, we merely need to transfer two types of messages between only two nodes. Even in the worst case, all types of messages just pass through part of the nodes in the new tree. Hence, the ISF approach may enjoy a significant reduction in computational cost compared with the naive method. To verify the theoretical analysis, we also conduct experiments to evaluate the performance of the ISF approach and the results in Fig. 11 demonstrate that this incremental approach can effectively enhance the efficiency of the search algorithm.

Further improvement: The performance of this algorithm is sometimes impaired by noisy non-free nodes. These noisy nodes cannot connect to other non-free nodes to generate complete answers, but they have very high importance values. In some cases, the branch and bound search algorithm wastes a lot of time to enumerate the candidate trees grown from these noisy nodes because they have higher upper bound estimates. To alleviate such problems and obtain more accurate upper bound estimates, we build indexes on the database graph, as discussed in the following section. The index improves the evaluation of the upper bounds of the candidate trees during their generation in Lines 12 and 16 of Algorithm 1.

5. Indexing techniques

We design indexing methods to further improve the search efficiency of the branch and bound search algorithm.

5.1. Naive indexing

In a naive indexing scheme, we could compute offline and store the shortest distances between any two nodes in the data graph, $DS(v_i, v_j)$, and the minimal loss of messages due to dampening when surfers move from one node to another, $LS(v_i, v_j)$. The shortest distance information gives more pruning power at run-time. When expanding a candidate tree, a non-free node can be ignored if it is too far away from the root of the candidate tree such that the diameter of the result

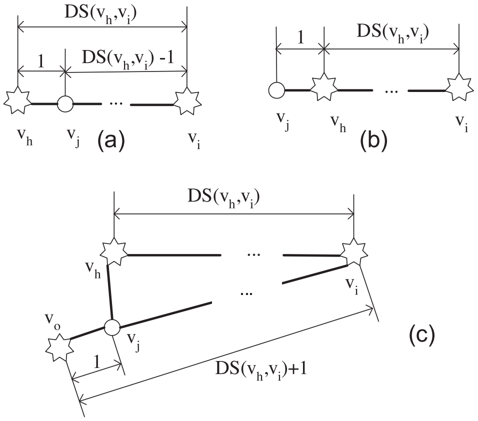

Fig. 7. The shortest distance between a star node and a non-star node.

tree would exceed the limit *D*. The information on the minimal loss of messages improves the accuracy of the upper bound estimate used in the search.

The main problem of this indexing scheme is that its space complexity is $O(|\mathbb{V}|^2)$, too big even for databases of moderate sizes. This leads us to the development of better indexing schemes. In particular, we consider the special case where the database has a star schema, which is common in OLAP applications, and introduce a new indexing method called *star indexing*. This indexing method dramatically reduces the space requirement and strikes a balance between the size of the index and its pruning power.

5.2. Star indexing

In star indexing, only part of the nodes, which we call star nodes, are indexed. All nodes in the star table are considered star nodes, where a star table is defined as a table such that if it is removed, all the remaining tuples in the database become disconnected. For example, in Fig. 1(b), the Movie table in the IMDB database is a star table. Similarly, the Paper table is a star table in the DBLP database as shown in Fig. 1(a). It is possible that tuples in the database cannot be completed disconnected by the removal of a single star table. In this case, we can have multiple star tables. The nodes in all of the star tables are indexed, maintaining the same information as in naive indexing. This design results in a significant reduction of the nodes that have to be indexed.

There are several different cases in leveraging the proposed star index to speed up the branch and bound algorithm by evaluating the relationship between two nodes.

Case 1: Both nodes are star nodes. The shortest distance $DS(v_i, v_j)$ and the minimal loss of the messages $LS(v_i, v_j)$ can be retrieved directly from the index.

Case 2: A star node v_i and a non-star node v_j. In this case, we first retrieve the set of star nodes directly connected to the v_j. Let it be $S_{star}(v_j)$. Then, the shortest distances and minimal message losses are retrieved for all $v \in S_{star}(v_j)$ and v_i. We select the star node v_h from $S_{star}(v_j)$ that gives the lowest values among all retrieved shortest distances. The shortest distance between v_h and v_i is denoted by $DS(v_h, v_i)$. If v_j is on the shorted path from v_h to v_i, the shortest distance between v_i and v_j is $DS(v_h, v_i) - 1$. This is shown in Fig. 7(a). The star nodes are represented by star symbols, and the other nodes are depicted as circles. Other possible values for the shortest distance between v_j and v_i are $DS(v_h, v_i) + 1$ and $DS(v_h, v_i)$, illustrated in Fig. 7(b) and Fig. 7(c) respectively. The latter value is given when there is another star node v_o that connects directly to v_j, which has a shortest distance $DS(v_h, v_i) + 1$ to node v_i, and has v_j on its shortest path. Therefore, we cannot get the accurate shortest distance in this case. In our search algorithm, we use the lower bound, $DS(v_h, v_i) - 1$ as the shortest distance. The pruning power is reduced because some nodes that should be pruned might be kept due to the inaccuracy thus introduced. This illustrates the trade-off between the index size and its pruning power. The minimal loss of messages is calculated in a similar way.

Case 3: Both nodes are non-star nodes. We expand both non-star nodes to two sets of star nodes and calculate the approximate shortest distance in a manner similar to the second case.

5.3. Group indexing

To achieve further reduction in index size, we extend star indexing and propose *group indexing*. Under this scheme, several star nodes can be grouped as one index unit. Only the inter-group relationships are indexed. A group is selected as follows. For a given group center and a user tunable parameter named group diameter *d*, all star nodes within distance *d* and not included in other groups are selected as members of this group. If a star node is chosen as a group center, the number of the star nodes in this group is called the *group capacity* of this center star node.

Table 1
Benefits of the model.

	Characteristics of the model
1	The number of the source messages is proportional to the importance of the source node. **Effect:** The important non-free nodes are favored.
2	The number of the messages is dampened when they pass through a node. **Effect:** Messages in a larger JTT have to go through more intermediate nodes before reaching other non-free nodes and thus tend to be dampened more heavily, resulting in lower scores. Smaller trees are preferred.
3	The relation between the dampening rate and the node importance is monotonic. **Effect:** Preference is given to important free nodes in connecting the non-free nodes.
4	The score of the tree is not dominated by the free nodes. **Effect:** The free node domination problem is avoided.

The group centers are decided by a greedy algorithm that aims to minimize the number of groups. This algorithm always selects the star node with the maximal group capacity as a group center in each iteration. The groups are mutually exclusive, i.e., a node cannot appear in more than one group. The shortest distances from every node in one group to every node in another group are evaluated; the lowest value among them is used as the shortest distance of the two groups and stored in the index. The minimal loss of the messages between the two groups is calculated and stored in a similar way. Apparently, the information stored in this manner is less precise than that in star indexing, but enjoys the benefit of being more space efficient. When the group diameter d is chosen to be a typical number 2, the number of centers is only 4.23% of the number of the star nodes for the IMDB database we experimented with.

Benefits: Star indexing and group indexing enhance the efficiency of the branch and bound search algorithm in two ways. First, it can provide a more accurate complete estimate. Without the index, if a supplement node is not directly connected to the root of the candidate tree, we just assume that it is connected to the best neighbor of the root. However, the supplement node may be far from the root, or even not connected at all. The minimal loss of messages between two nodes, which we can get from the index, is used to improve the accuracy of the complete estimate. Second, we can avoid computing the potential estimate and get a more accurate evaluation of the contribution from the additional nodes. The contributions of all non-free nodes that are not in the candidate tree can be calculated using the index. A greedy algorithm retrieves them in descending order of the contributions, and adds them into the candidate tree until the score of the result tree starts to decrease. Because the star and group indexes do not store the precise values for the loss of the messages, the calculated contribution is not an exact value but rather an upper bound.

6. Experiments

6.1. Experiment setting

To evaluate the effectiveness and efficiency of CI-Rank, we conduct extensive experiments on two large-scale data sets and a real user query log.

The data sets used in the experiments are the Internet Movie Database (IMDB),[1] and the DBLP data.[2] Both data sets are modeled as graphs. The IMDB data contains 3,378,743 nodes and 28,482,926 edges, and the DBLP data contains 2,132,821 nodes and 8,446,804 edges respectively. The schemas of the two databases are shown in Fig. 1.

Similar to the previous study [10], the edge weights are chosen experimentally when modeling the databases as graphs in our experiments. We find the edge weights listed in Table 2 work well in our experiments. These weights will be normalized in the graph. For example, one movie has three out edges, which point to an actor, a director and a producer. Before the normalization, the weights of these edges are 1.0, 1.0, and 0.5. The random walk model requires that the weights of out edges of a node sum to 1.0. Therefore, the weights of above three out edges are normalized to 0.4, 0.4, and 0.2. There is another problem in modeling the IMDB data set due to its normalization. The same person may have several copies because he/she has different roles. For example, Mel Gibson is the director of the movie Braveheart (1995). At the same time, he is an actor in it. When we map the IMDB data set to a graph, two nodes, director Mel Gibson and actor Mel Gibson, are created. As a result, the importance value belonging to the same person is split and distributed into two different nodes. To avoid this case, we merge such nodes into a single node. In the above example, we only have one Mel Gibson node in the graph. This node has two different edges, directing and acting, that point to the node of the movie Braveheart (1995) (Table 1).

A real user query log is used in our experiments. In 2006, AOL provided the user query log accumulated in their search engine for three months. The log, which can be obtained from a mirror site,[3] contains 20 million queries issued by 650,000 users. Each record in this log has both the query content and the URL clicked by the user. Only the records whose clicked URL contains http://www.imdb.com are collected. The total number of such records is 81,250. Among them, we manually

---

[1] http://www.imdb.com/interfaces.

[2] http://dblp.uni-trier.de/xml/.

[3] http://www.gregsadetsky.com/aol-data/.

Table 2
The edge weights.

Data set	Edge type	Weight
IMDB	Actor → Movie	1.0
	Movie → Actor	1.0
	Actress → Movie	1.0
	Movie → Actress	1.0
	Director → Movie	1.0
	Movie → Director	1.0
	Producer → Movie	0.5
	Movie → Producer	0.5
	Company → Movie	0.5
	Movie → Company	0.5
DBLP	Conference → Paper	0.5
	Paper → conference	0.5
	Author → Paper	1.0
	Paper → Author	1.0
	Citing paper → Cited paper	0.5
	Cited paper → Citing paper	0.1

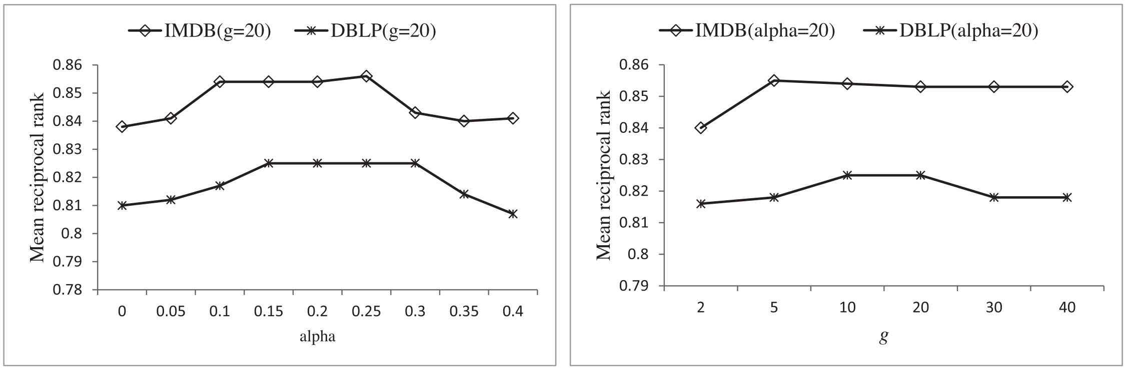

(a) Effect of α (b) Effect of *g*

Fig. 8. The effect of parameters on the mean reciprocal rank.

labeled 29,078 frequent queries. For robustness, each of those queries must have appeared at least three times in the log. These labeled queries are used as user feedback to bias the CI-Rank model. 44 complex queries are selected from the user query log to verify the effectiveness and efficiency. These complex queries, which match at least two nodes, have clear meaning and no ambiguity in the manual labeling. Since the AOL log does not contain any queries related to DBLP, 20 synthetic queries are used instead. The method to generate the synthetic queries is explained in the following subsection.

CI-Rank is implemented in Java. The index is built using Apache Lucene. All the experiments are run on an IBM Linux server with a 3.0GHz Intel Dual Core processor, 4GB of RAM, and 2TB SATA HD RAID.

In the experiments, we implement SPARK's [7] scoring function on the database graph, as well as BANKS [2], to compare their performance with that of CI-Rank. We also design a ranking approach called PageRankSearch as another baseline. PageRankSearch adapts PageRank [14] to keyword search over relational databases. In this approach, we first view the data graph representing the database as a graph of webpages (in the original PageRank context), where each node corresponds to a page and the edge represents a reference from one page to another, and then use a simplified algorithm [15] to compute the PageRank value of each node. PageRankSearch defines the weight of a Steiner tree as the sum of the PageRank values of all nodes in this tree, and the Steiner trees can be ranked according to their weights in a descending order. Since the weight of a Steiner tree includes the PageRank values of both free nodes and non-free nodes, the main drawback of PageRankSearch is that the free nodes with large PageRank values will have excessive impact on the results which may lead to degradation in the effectiveness of ranking. This will be empirically evaluated in the Section 6.2 (Fig. 8).

6.2. *Effectiveness*

We use two metrics, mean reciprocal rank and precision, to measure the effectiveness. The reciprocal rank is the inverse of the rank of the best answer. The best answer is decided by a user study. Five graduate students were invited to select their most preferred answer for each query. For a query, the best answer is decided by majority voting. In the case of a tie, all of the answers are considered the best. Another measure, precision, is the fraction of the answers generated that

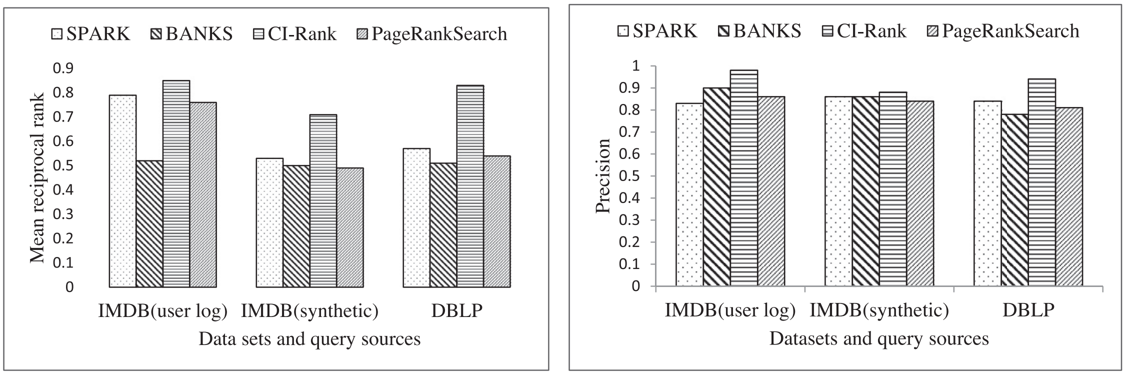

(a) Mean reciprocal rank (b) Precision

Fig. 9. Effectiveness comparison of CI-Rank with SPARK, BANKS, and PageRankSearch.

are relevant to the query. Graded relevance levels are used. If a relevant answer does not contain all input keywords, its relevance level is penalized by the percentage of the missed keywords.

We first study the effect of the two parameters in CI-Rank, namely, α and g. Both parameters defined in Section 3.3 are used to control the message dampening process inside a node. Recall that α is the probability a message carrier is willing to take the messages in one talk. g is the clique size of the talk. The importance values of the graph nodes are calculated by the iterations in Eq. (1). A typical value, 0.15, is used for the teleportation constant c in this equation.

Fig. 8(a) and (b) illustrate how the effectiveness is affected by the parameters α and g. The change of the mean reciprocal rank shows the same pattern for both data sets. The result is better when α is between 0.1 and 0.25 than when the α is out of this range. Similarly, the parameter g leads to better accuracy when it ranges from 10 to 20. We choose $\alpha = 0.15$ and $g = 20$ for the following experiments since CI-Rank has the best performance with them for both IMDB and DBLP data sets. The value range of the dampening rate in CI-Rank is controlled by α and g. α is the minimal value of the dampening rate. If α is fixed, the maximal value of the dampening rate decreases when g increases.

To compare CI-Rank and existing methods, we introduce the ranking strategies of SPARK, BANKS, and PageRankSearch. The result of CI-Rank is obtained with the default parameters $\alpha = 0.15$ and $g = 20$. No user feedback is used. Two different query sets are used for IMDB. One is obtained from the AOL user log, and the other synthetically generated. Interestingly, most of the complex queries obtained from the AOL user log can be matched by two directly connected nodes, which is typical of Web search. Considering that a database search user is more likely to retrieve the information that links two or more separated elements, a synthetic query set is also generated for IMDB. It has the same number of queries, 20, as the DBLP query set. Both synthetic query sets follow the same pattern. 50% of the queries are matched by two non-free nodes that are not directly connected, and 20% of the queries cover three or more non-free nodes. The remaining queries can be matched by either a single node or two directly connected nodes.

The comparison in terms of mean reciprocal rank is shown in Fig. 9(a). For queries obtained from the AOL user log, the effectiveness of CI-Rank, SPARK, and PageRankSearch is close (0.85, 0.79 and 0.76 respectively), all better than BANKS. For synthetic queries on both IMDB and DBLP, the mean reciprocal ranks of SPARK, BANKS and PageRankSearch are about 0.5, much lower than that of CI-Rank. The reason is that for most user log queries, the answers are directly connected non-free nodes, with no free nodes required. Only 11.4% of them require free nodes to connect the matched non-free nodes. For synthetic DBLP queries, this percentage is chosen to be 50%. This validates the effectiveness of our ranking function in selecting the free nodes to connect non-free nodes, as well as the flaw of PageRankSearch that the free nodes excessively impact the ranking results.

Fig. 9(b) shows the comparison of the methods in terms of precision. The precision of CI-Rank is higher than 0.9 in three different experiments. SPARK, BANKS and PageRankSearch also have high precision values, which are more than 0.85 for IMDB and 0.75 for DBLP. The difference in precision between CI-Rank and SPARK can be primarily attributed to those long queries that match three or more non-free nodes.

6.3. User feedback

Next we evaluate the effect of user feedback. This experiment is only performed on the IMDB data set since we do not have the real user query log for the DBLP data set. In the unbiased model, the teleportation vector **u** in Eq. (1) has a uniform distribution. Now it is populated using the query log as follows. If a node v_i shows up k times in the query result chosen by the user according to the log, its teleportation data u_i is set to be proportional to k. Under this rule, for a node v_i that does not appear in the log, its u_i should be set to zero. However, considering that the log does not span a long time and that we have only labeled the frequent queries, it may penalize such a node too much if u_i is set to zero. Instead, we assign it a small value by assuming that the unlabeled queries are uniformly distributed over the nodes that are not covered by

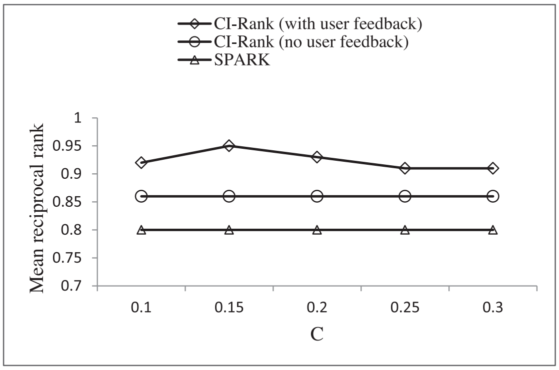

Fig. 10. The effect of user feedback on mean reciprocal rank for IMDB.

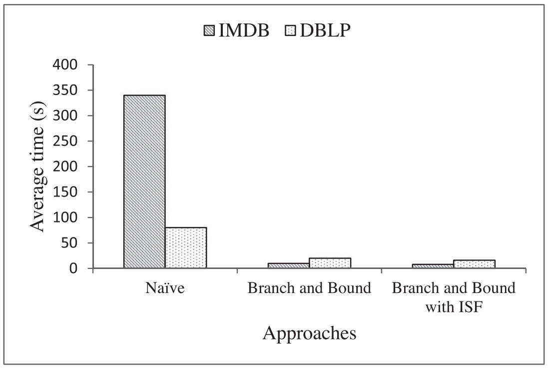

Fig. 11. The efficiency of the naive algorithm and the branch and bound algorithm.

the log. This small value is $1.9e-7$ in our experiment. For comparison, the minimum u_i for the labeled nodes is $3.7e-5$. Fig. 10 shows the effect of incorporating the user feedback in terms of mean reciprocal rank. The horizontal axis in this figure corresponds to the teleportation constant c in Eq. (1). By varying this parameter, we adjust the bias strength of the model. Clearly, CI-Rank with user feedback gives the best result, achieving a mean reciprocal rank of 0.95 when c is 0.15. Over all c values tested, the mean reciprocal rank is always higher than 0.9 over all c values tested, averaging a 5%-10% improvement over CI-Rankwithout user feedback. It indicates that incorporating user feedback does effectively improve the quality of ranking.

6.4. Efficiency

We first show the efficiency comparison between the naive algorithm and the branch and bound algorithm. Because the naive algorithm can easily run out of memory when the whole datasets are used in their entirety, we obtain uniformly samples of the original datasets, with the size of each being 10% of the original. The experiments are performed on the sample datasets, and the results are presented in Fig. 11. It is evident that the branch and bound algorithm significantly outperforms the naive one.

The evaluation of star index and group index is performed on the full datasets. Here, we set the group diameter as 2. The search time for the answer tree is measured for different maximal tree diameters D, and the results are presented in Fig. 12(a) and (b). Each result is the average search time for the top-5 answers for the IMDB and DBLP queries. It is evident from these two graphs that the indexes reduces the search time considerably on both datasets, and star index performs better than group index. When the maximal tree diameter D decreases, the search time generally drops. For all maximal tree diameters, the average search time is less than 10 seconds when the star index is used. Note that when the maximal tree diameter D is set to a small value (e.g., 2 or 3), the number of the full answer trees for some queries is less than 5, the value of k in top-k search.

7. Related work

Keyword search has been extensively studied in the literature [1,4–7,10,19–21]. See [22] for an excellent survey. We have examined a few representative approach in Section 2. Here we briefly mention a couple other studies that are related to our work.

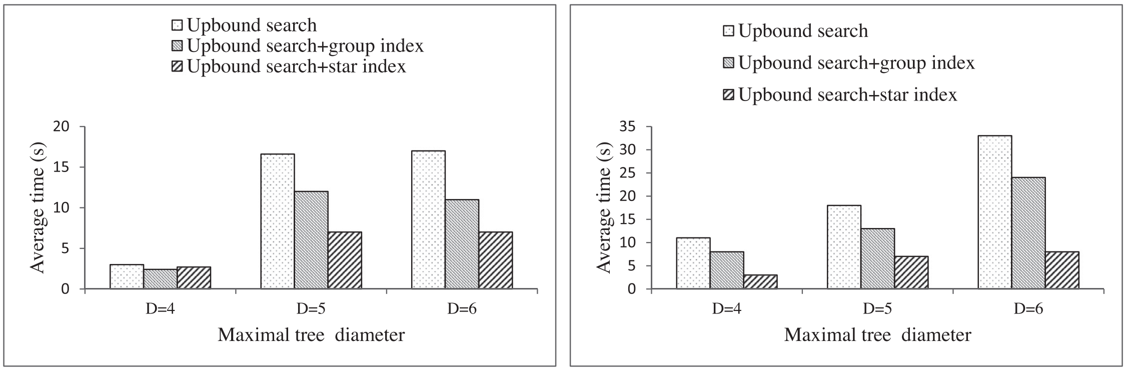

(a) IMDB (b) DBLP

Fig. 12. The average search time for queries on two datasets.

Instead of generating answer tree, EASE [9] looks for the most relevant subgraph, constrained by a maximum radius specified by the user. The ranking function in EASE is still IR-style. Thus, it does not consider the node importance.

CEPS [23] uses random walk with restarts (RWR), a variant of the random walk model, to find the best center-piece subgraph for a given graph and several pre-selected query nodes. The result of this algorithm is a subgraph rather than a tree, which is different from the query answer form used in our studies. A problem with this approach is that the query nodes are pre-selected in this model. For keyword search in databases, every input keyword will hit a lot of non-free tuples. It is unclear how to best choose candidates from all non-free tuples as the query nodes.

MeanKS [24] studies meaningful keyword search in relational databases with complex schema. This system allows the user to specify the “role” of each keyword through a user interface, and returns the top minimal join trees that not only contain all the query keyword roles but also reflect the user interests. The ranking of minimal join trees is based on the importance of nodes, edges, or combinations of them.

Bergamaschi et al. [25] propose QUEST, a keyword search system over relational databases implementing a three-step schema-based approach. It first determines how the keywords in the query correspond to the structural elements of the database using Hidden Markov Models, and then identifies the structure of the queries that can be generated from a given configuration. Finally, it decides on the combination of keyword mappings into data structures and paths connecting these data structures that most satisfy the requirements of users.

Coffman et al. [26] present a thorough empirical performance evaluation of relational keyword search systems.

8. Conclusions and future work

We have proposed CI-Rank, a new approach for keyword search in databases. CI-Rank considers both the importance of individual nodes and the cohesiveness of the answer tree in a balanced way, and avoids a lot of the pitfalls of the ranking functions employed in previous methods. We presented a branch and bound algorithm to support the efficient generation of top-k query answers. The efficiency of the algorithm is further improved by star indexing. All solutions are implemented and evaluated against the IMDB and DBLP datasets. Real user query logs are employed in the evaluation.

In future work, we would like to study user preference adaptation. One possible direction is to consider how to improve the model such that user feedback can be used to adjust not only the importance values of the nodes, but also the weights of the edges of the database graph. Another direction is to study how to combine the importance-based ranking used in our approach and IR-style ranking to produce better ranking functions.

Acknowledgment

This work was supported in part by the 973 Program (Grant No. 2015CB352502), the National Natural Science Foundation of China (Grants No. 61272092, No. 61572289), the Shandong Provincial Natural Science Foundation Grants (Grants No. ZR2015FM002, ZR2016FB14), the Science and Technology Development Program of Shandong Province (Grant No. 2014GGE27178), and NSERC Discovery Grants.

Appendix A. Proof of Lemma 1

Proof. Suppose we have an arbitrary full answer tree T grown from the candidate tree C. Now we prove that $score(T) <= ub(C) = \max(ce(C), pe(C))$.

(0) We divide the non-free nodes in T into three groups. *Group* a consists of the non-free nodes in C, denoted as $v_1, v_2, \ldots, v_n$. Since C is not a complete tree, the keywords that are not covered by C are matched by some other non-free nodes in T. We choose an arbitrary combination (*group* b) of these nodes that just covers the missed keywords in C, denoted as $w_1, w_2, \ldots, w_m$. The remaining non-free nodes form *group* c, and are denoted as $u_1, u_2, \ldots, u_l$.

(1) We first prove that the average node score of *groups* a and b is less than or equal to the complete estimate, i.e., $(\sum_{i=1}^{n} score(v_i) + \sum_{i=1}^{m} score(w_i))/(n+m) \le ce(C)$. The steps are as follows. In the hypothetical tree that is used to evaluate $ce(C)$, the best complement nodes are attached directly to the root of C. Denote them as $w'_1, w'_2, \ldots, w'_m$. Because they are the best complement nodes, the importance of w'_i is higher than or equal to that of w_i. For any non-free node v_i in C, $score(v_i)$ is its node score in tree T. Let $score'(v_i)$ denote its node score in the hypothetical tree. Since $w'_1, w'_2, \ldots, w'_m$ have higher importance values and connect more closely than $w_1, w_2, \ldots, w_m$, v_i will receive more messages from $w'_1, w'_2, \ldots, w'_m$. If the non-free nodes from *group c* are not considered, $score(v_i) \le score'(v_i)$. Because the node score is determined by the least populous type of messages, the additional non-free nodes from *group c* can only make v_i have the same or lower node score. Consequently, $score(v_i) \le score'(v_i)$ still holds when *group c* is considered. Next we compare the score of $score(w_j)$ and $score'(w'_j)$. Messages from v_i are less dampened before they are accepted in w'_j because w'_j has higher importance value and closer connection to v_i than w_j. The direct connection of w'_j to another *group* b node w'_k is considered in the hypothetical tree as well. Thus, the number of messages received from w'_k in w'_j is greater than or equal to that from w_k in w_j. Consequently, w'_j receives more messages than w_j for any message type. According to the definition of node score, $score(w_j) \le score'(w'_j)$. Combining the result of the above two steps leads to

$$\left[\sum_{i=1}^{n} score(v_i) + \sum_{i=1}^{m} score(w_i)\right] / (n+m) \tag{5}$$

$$\le \left[\sum_{i=1}^{n} score'(v_i) + \sum_{i=1}^{m} score'(w'_i)\right] / (n+m) = ce(C) \tag{6}$$

(2) We then prove that the average node score of *group* c is less than or equal to the potential estimate, $(\sum_{i=1}^{l} score(u_i))/l \le pe(C)$. The potential estimate is chosen as the max node score an additional non-free node can achieve. Thus, $pe(C) \ge s(u_i)$. Then, we can get

$$\sum_{i=1}^{l} score(u_i)/l \le \sum_{i=1}^{l} pe(C)/l = pe(C) * l/l = pe(C) \tag{7}$$

(3) According to the definition, the score of the tree T is

$$score(T) = \frac{\sum_{i=1}^{n} score(v_i) + \sum_{i=1}^{m} score(w_i) + \sum_{i=1}^{l} score(u_i)}{n+m+l} \tag{8}$$

Applying Eqs. (6) and (7) into Eq. (8) results in

$$\begin{aligned} score(T) &\le \frac{(n+m) * ce(C) + l * pe(C)}{n+m+l} \\ &\le \frac{(n+m+l) * max(ce(C), pe(C))}{n+m+l} \\ &= max(ce(C), pe(C)) = ub(C) \qquad \square \end{aligned}$$

Appendix B. Proof of Lemma 2

Proof. We set $ce(C(r_0))$and $pe(C(r_0))$ as the complete estimate and potential estimate of $C(r_0)$ respectively, and $ce(C(r_g))$ and $pe(C(r_g))$ are the complete estimate and potential estimate of $C(r_g)$ respectively. To prove Lemma 2, we also use the groups a and b as defined in Appendix A, i.e., $C(r_0)$ contains the non-free nodes $v_1, v_2, \ldots, v_n$. Suppose that a full tree T can be extended from $C(r_0)$ by adding the non-free nodes $w_1, w_2, \ldots, w_m$, and that $w'_1, w'_2, \ldots, w'_m$ are in the hypothetical tree. We redefine the group c so that it contains the non-free nodes $u_1, u_2, \ldots, u_k$ that can be connected to T to maximize $score(T)$. According to the above definitions, we can get

$$ce(C(r_0)) = \sum_{i=1}^{n} score(v_i) + \sum_{i=1}^{m} score(w'_i)$$

$$ce(C(r_g)) = \sum_{i=1}^{n} score(v_i) + score(w_j) + \sum_{i=1}^{m} score(w'_i),\ j \ne i$$

where node w_j represents the node r_g. Because $score(w_j) \leq score(w'_j)$, $ce(C(r_g)) \leq ce(C(r_0))$. Similarly, we can obtain

$$pe(C(r_0)) = score(T) + \sum_{i=1}^{k} score(u'_i)$$

$$pe(C(r_g)) = score(T) + score(u'_j) + \sum_{i=1}^{k} score(u'_i), i \neq j$$

Therefore, $pe(C(r_g)) \leq pe(C(r_0))$. Hence, $ub(C(r_g)) \leq ub(C(r_0))$. □

References

[1] S. Agrawal, S. Chaudhuri, G. Das, DBXplorer: A system for keyword-based search over relational databases, in: ICDE, IEEE Computer Society, Washington, DC, USA, 2002, p. 5.
[2] G. Bhalotia, A. Hulgeri, C. Nakhe, Chakrabarti, Keyword searching and browsing in databases using BANKS, in: ICDE, IEEE Computer Society, Washington, DC, USA, 2002, p. 431.
[3] V. Kacholia, S. Pandit, S. Chakrabarti, S. Sudarshan, R. Desai, H. Karambelkar, Bidirectional expansion for keyword search on graph databases, in: VLDB, VLDB Endowment, 2005, pp. 505–516.
[4] V. Hristidis, Y. Papakonstantinou, DISCOVER: keyword search in relational databases, in: VLDB, VLDB Endowment, 2002, pp. 670–681.
[5] V. Hristidis, L. Gravano, Y. Papakonstantinou, Efficient IR-style keyword search over relational databases, in: VLDB, VLDB Endowment, 2003, pp. 850–861.
[6] F. Liu, C. Yu, W. Meng, A. Chowdhury, Effective keyword search in relational databases, in: SIGMOD, ACM, New York, NY, USA, 2006, pp. 563–574. http://doi.acm.org/10.1145/1142473.1142536.
[7] Y. Luo, X. Lin, W. Wang, X. Zhou, SPARK: top-k keyword query in relational databases, in: SIGMOD, ACM, New York, NY, USA, 2007, pp. 115–126. http://doi.acm.org/10.1145/1247480.1247495.
[8] X. Yu, H. Shi, Ci-rank: Ranking keyword search results based on collective importance, ICDE, 2012.
[9] G. Li, B.C. Ooi, J. Feng, J. Wang, L. Zhou, EASE: an effective 3-in-1 keyword search method for unstructured, semi-structured and structured data, in: SIGMOD, ACM, New York, NY, USA, 2008, pp. 903–914. http://doi.acm.org/10.1145/1376616.1376706.
[10] A. Balmin, V. Hristidis, Y. Papakonstantinou, ObjectRank: authority-based keyword search in databases, in: VLDB, VLDB Endowment, 2004, pp. 564–575.
[11] V. Hristidis, H. Hwang, Y. Papakonstantinou, Authority-based keyword search in databases, ACM Trans. Database Syst. 33 (1) (2008).
[12] S. Brin, L. Page, The anatomy of a large-scale hypertextual web search engine, Comput. Netw. ISDN Syst. 30 (1–7) (1998) 107–117. http://dx.doi.org/10.1016/S0169-7552(98)00110-X.
[13] L. Page, S. Brin, R. Motwani, T. Winograd, The PageRank Citation Ranking: Bringing Order to the Web., Technical Report, Stanford InfoLab, 1999.
[14] L. Page, S. Brin, R. Motwani, T. Winograd, The pagerank citation ranking: bringing order to the web. (1999).
[15] Wikipedia, Simplified algorithm, 2016, (https://en.wikipedia.org/wiki/PageRank/).
[16] K.-H. Vik, P. Halvorsen, C. Griwodz, Evaluating steiner-tree heuristics and diameter variations for application layer multicast, Comput. Netw. 52 (15) (2008) 2872–2893. http://dx.doi.org/10.1016/j.comnet.2008.06.003.
[17] B. Ding, J.X. Yu, S. Wang, L. Qin, X. Zhang, X. Lin, Finding top-k min-cost connected trees in databases, in: ICDE, 2007, pp. 836–845.
[18] J.A. Rodger, Neurois Knowledge Discovery Approach to Prediction of Traumatic Brain Injury Survival Rates: A Semantic Data Analysis Regression Feasibility Study, in: Information Systems and Neuroscience, Springer, 2015, pp. 1–8.
[19] L. Qin, J. Yu, L. Chang, Ten thousand sqls: parallel keyword queries computing, PVLDB 3 (1) (2010) 58–69.
[20] A. Baid, I. Rae, J. Li, A. Doan, J.F. Naughton, Toward scalable keyword search over relational data, PVLDB 3 (1) (2010) 140–149.
[21] D. Xin, Y. He, V. Ganti, Keyword++: a framework to improve keyword search over entity databases, PVLDB 3 (1) (2010) 711–722.
[22] J.X. Yu, L. Qin, L. Chang, Keyword Search in Databases, Morgan and Claypool Publishers, 2010.
[23] H. Tong, C. Faloutsos, Center-piece subgraphs: problem definition and fast solutions, in: KDD, ACM, New York, NY, USA, 2006, pp. 404–413. http://doi.acm.org/10.1145/1150402.1150448.
[24] M. Kargar, A. An, N. Cercone, P. Godfrey, J. Szlichta, X. Yu, Meanks: meaningful keyword search in relational databases with complex schema, in: SIGMOD, ACM, 2014, pp. 905–908.
[25] S. Bergamaschi, F. Guerra, M. Interlandi, R. Trillo-Lado, Y. Velegrakis, Quest: a keyword search system for relational data based on semantic and machine learning techniques, in: VLDB, VLDB Endowment, 2013, pp. 1222–1225.
[26] J. Coffman, A.C. Weaver, An empirical performance evaluation of relational keyword search techniques, TKDE 26 (1) (2014) 30–42.