

Federation University ResearchOnline

<https://researchonline.federation.edu.au>

Copyright Notice

This is the post-peer-review, pre-copyedit version of an article published in World Wide Web. The final authenticated version is available online at:

<https://doi.org/10.1007/s11280-017-0503-8>

Copyright © 2017, Springer Science+Business Media, LLC

Use of the Accepted Manuscript is subject to [AM terms of use](#), which permit users to view, print, copy, download and text and data-mine the content, for the purposes of academic research, subject always to the full conditions of use. Under no circumstances may the AM be shared or distributed under a Creative Commons, or other form of open access license, nor may it be reformatted or enhanced.

See this record in Federation ResearchOnline at:

<http://researchonline.federation.edu.au/vital/access/HandleResolver/1959.17/182317>

No-But-Semantic-Match: Computing Semantically Matched XML Keyword Search Results

Mehdi Naseriparsa · Md. Saiful Islam · Chengfei Liu · Irene Moser

Received: date / Accepted: date

Abstract Users are rarely familiar with the content of a data source they are querying, and therefore cannot avoid using keywords that do not exist in the data source. Traditional systems may respond with an empty result, causing dissatisfaction, while the data source in effect holds semantically related content. In this paper we study this no-but-semantic-match problem on XML keyword search and propose a solution which enables us to present the top- k semantically related results to the user. Our solution involves two steps: (a) extracting semantically related candidate queries from the original query and (b) processing candidate queries and retrieving the top- k semantically related results. Candidate queries are generated by replacement of non-mapped keywords with candidate keywords obtained from an ontological knowledge base. Candidate results are scored using their cohesiveness and their similarity to the original query. Since the number of queries to process can be large, with each result having to be analyzed, we propose pruning techniques to retrieve the top- k results efficiently. We develop two query processing algorithms based on our pruning techniques. Further, we exploit a property of the candidate queries to propose a technique for processing multiple queries in batch, which improves the performance substantially. Extensive experiments on two real datasets verify the effectiveness and efficiency of the proposed approaches.

Keywords XML Keyword Query · No-Match · Semantics · Ontology

M. Naseriparsa · Md. Saiful Islam · C. Liu · I. Moser
E-mail: {mnaseriparsa,mdsaifulislam,cliu,imoser}@swin.edu.au
Swinburne University of Technology

1 Introduction

Users who query data sources using keyword searches often are not familiar with the data source schema or the appropriate query language. For the query to succeed, the keywords have to have matches in the data source. Failing this, an empty result is returned even when semantically related content exists. When keywords have indirect mappings in a data source that cannot be found by traditional systems, the user faces the no-but-semantic-match problem.

Example 1 . Consider a user submitted a keyword query $q_0 = \{Jack, lecturer, class\}$ on XML database given in Fig. 1 and would like to find information about the professor *Jack*. Using conjunctive keyword search, traditional systems will show an empty result because there is no occurrence for the keywords *lecturer* and *class* in the data source. However, the keyword *lecturer* has a semantic connection to *academic* and *full professor* while the keyword *class* is semantically related to *course*, *grade* and *event* which exist in the data source and could generate results that might interest the user.

The XML keyword search has been addressed by researchers before. The concept of Lowest Common Ancestor (LCA) was first proposed by Guo et al. [13] to extract XML nodes which contain all query keywords within the same subtree. Xu and Papakonstantinou [26] introduced the concept of Smallest Lowest Common Ancestor (SLCA) to reduce the query result to the smallest tree that contains all keywords. Sun, Chan and Goenka [23] extended this work by applying the SLCA principle to logical OR searches. Hristidis et al. [14] explored the trees below LCA to provide information about the proximity of the keywords in the document. None of the existing studies use the SLCA semantics

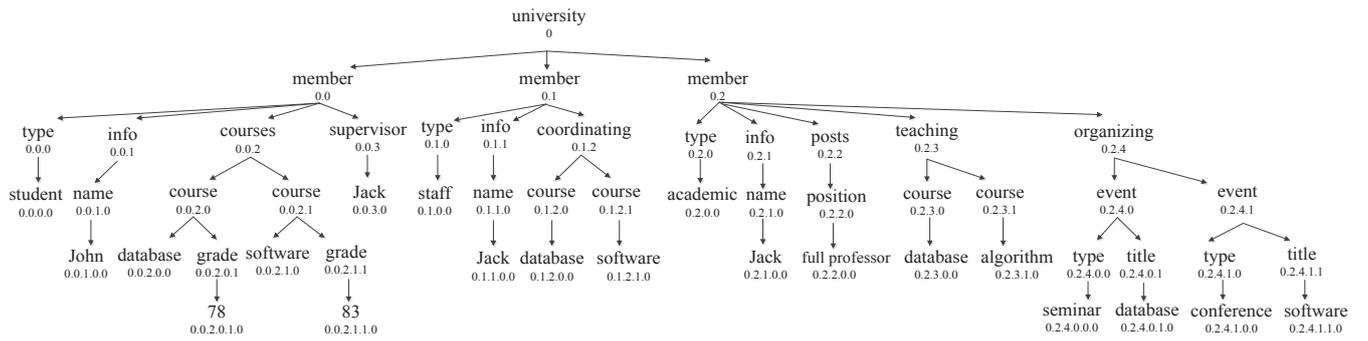


Fig. 1: A part of XML data

to provide a solution when one or more keywords do not exist in the database. In this paper, we adapt the widely-accepted SLCA semantics and algorithms to retrieve meaningful results when some non-mapped keywords are submitted to the system.

When a query encounters the no-but-semantic-match problem, we need to find candidate keywords for the non-mapped keywords to produce non-empty results. Even though the non-mapped keywords may be semantically close to some items in data source, traditional systems do not attempt to discover them. To produce an answer to the user’s initial query, the candidate keywords must be semantically close to the non-mapped keywords. One way of fulfilling this requirement is to find substitutes for non-mapped keywords in an ontological knowledge base. Clearly, only candidate keywords that have a mapping in the data source can be selected as substitutes for a new query. Replacing each of the non-mapped keywords with one or more semantically related words that are known to exist in the database leads to a list of candidate queries. Depending on the number of available keywords, the number of potential queries and results can be impractically large. Hence the degree of semantic similarity with the original query is calculated for each candidate query before it is executed. Before the results can be presented to the user, results of poor quality in terms of cohesiveness must be eliminated to ensure all results are meaningful answers to the original query. Thus, to solve the no-but-semantic-match problem, two aspects are considered: (a) query similarity; and (b) result cohesiveness.

Example 2. Consider the keyword query $q_0 = \{Jack, lecturer, class\}$ presented in Example 1 on the database shown in Fig. 1. Keywords *lecturer* and *class* do not have a mapping in the data source and the traditional system generates an empty result for it. The ontological knowledge base [18] has 44 semantic counterparts for *lecturer* and 39 for *class*. All possible substitutions and their combinations are considered. In the extreme

case when all candidate keywords are available in the data source, $44 \times 39 = 1716$ queries are generated and each query may have several answers that have to be considered. When a high number of keywords have to be replaced and these keywords have many semantic counterparts, we may face an unmanageably large number of combinations that have to be analyzed for semantic similarity with the original query. Hence, there is a need to identify and remove less promising candidate queries early.

In this paper, we present a novel two-step solution to the no-but-semantic-match problem in XML keyword search. In the first step, semantically related candidate queries are created by replacing non-mapped keywords in the original queries with semantic counterparts and in the second step, the queries are processed and the top-k semantically related results retrieved. In order to present the top-k results to the user for evaluation, each result retrieved from the queries is separately analyzed in terms of its similarity to the original query and its cohesiveness in data source. Since there may be a large number of semantically related results, retrieving the top-k results is potentially costly. Therefore, we propose two pruning techniques, inter-query and intra-query pruning. Since the candidate queries are generated by replacing non-mapped keywords, some keywords are shared between the candidate queries. We exploit this property to propose a more efficient batch query processing technique to improve the performance substantially. The issue of finding semantically related results for queries with no-but-semantic-match problem has not been addressed in the context of semi-structured data before. Our contributions are as follows:

1. We are the first to formulate the no-but-semantic-match problem in XML keyword search.
2. We propose two pruning methods and an efficient approach of processing the no-but-semantic-match query.

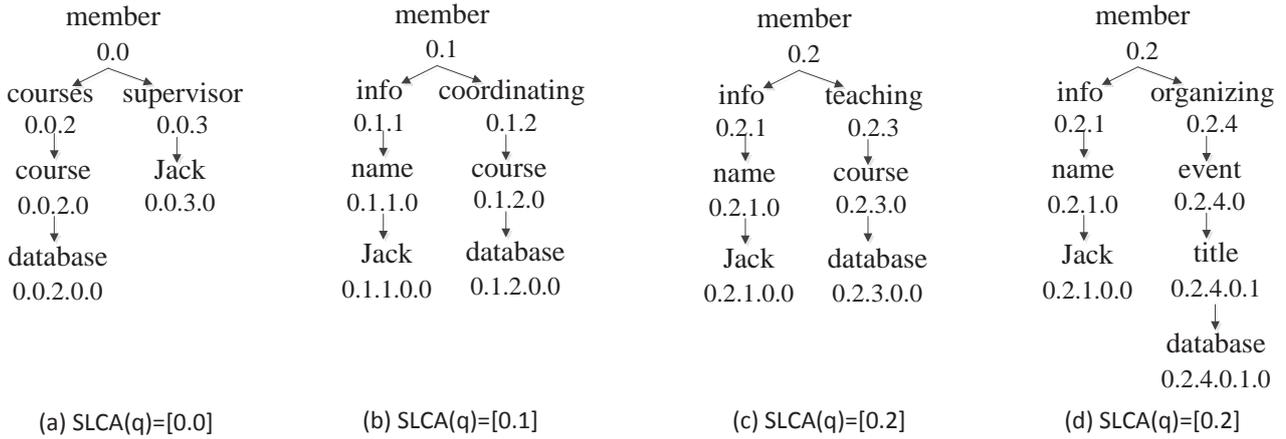


Fig. 2: SLCA subtree results for query $q = \{Jack, database\}$ executed on data given in Fig. 1

- Based on keywords the candidate queries have in common, we also propose a method to process multiple queries in a batch which improves the performance substantially.
- We conduct extensive experiments which verify the effectiveness and efficiency of our solutions on two real datasets.

The rest of the paper is organized as follows: Section 2 discusses XML keyword search and presents the no-but-semantic-match problem. Section 3 presents the details of our pruning ideas and the efficient processing of the no-but-semantic-match query. Section 4 presents the batch query processing scheme to further improve the performance. The experiments are presented in Section 5. Section 6 reviews the related work. Finally, Section 7 concludes our paper.

2 Background

2.1 Preliminaries

An XML document is an ordered tree T with labeled nodes and a designated root. All XML elements are treated as nodes containing information in T . There are parent-child and sibling relationships between the nodes. The depth of the tree is denoted as d , and the root node has a depth of 1. Each node v in the tree T is marked with a unique identifier in Dewey code, which describes the path from the root to the node v as a sequence of numbers separated by a dot (“.”). Sibling nodes have Dewey codes of equal length with a unique last number.

Example 3. Fig. 1 shows an XML tree which contains information about staff and students of a university. The root node’s Dewey code is 0. Dewey code 0.0.1

refers to a node containing information about a member of the university and the code prefix 0.0 refers to its parent node.

Keyword Match Node: A node m in the tree T is a match node for keyword k_i if it contains k_i . e.g., the match nodes for keyword $k_1 = database$ presented in Fig. 1 are: $m_1^1 = [0.0.2.0.0]$, $m_2^1 = [0.1.2.0.0]$, $m_3^1 = [0.2.3.0.0]$, and $m_4^1 = [0.2.4.0.1.0]$.

Keyword Inverted List: Each keyword k_i corresponds to a list S_i of entries and each entry corresponds to a node m which contains k_i in the tree T . e.g., the keyword inverted list for keyword $k_1 = database$ is $S_1 = \{[0.0.2.0.0], [0.1.2.0.0], [0.2.3.0.0], [0.2.4.0.1.0]\}$.

Smallest Lowest Common Ancestor (SLCA): Let $lca(m_1, \dots, m_n)$ returns the lowest common ancestor (LCA) of match nodes m_1, \dots, m_n . Then LCAs of query q on T are defined as $LCA(q) = \{v | v = lca(m_1, \dots, m_n), m_i \in S_i (1 \leq i \leq n)\}$. SLCA is a subset of LCAs which do not have other LCAs as child nodes and defined as $SLCA(q)$.

Example 4. In Fig. 2, for a keyword query $q = \{Jack, database\}$, there are 4 LCA nodes which are computed as: $LCA(q) = \{lca([0.0.3.0], [0.0.2.0.0]), lca([0.0.3.0], [0.1.2.0.0]), lca([0.1.1.0.0], [0.1.2.0.0]), lca([0.2.1.0.0], [0.2.3.0.0])\} = \{[0], [0], [0.1], [0.2]\}$. Since the LCA node $[0]$ is the ancestor node of $[0.0], [0.1]$ and $[0.2]$, it is not an SLCA and should be removed. Therefore, $SLCA(q) = \{[0.0], [0.1], [0.2]\}$.

Keyword Query and Subtree Result: In XML data, a keyword query q consists of a set of keywords $\{k_1, k_2, \dots, k_n\}$. A result $r = (v_{slca}, \{m_1, m_2, \dots, m_n\})$ for q is a subtree in T which contains all keywords $k_i \in q$. Here, we consider v_{slca} , the root of the subtree, an SLCA node, i.e. $v_{slca} \in SLCA(q)$.

Table 1: The list of symbols

Symbol	Meaning
$\lambda(q_0, q')$	Similarity score of q_0 to q'
α	Tuning parameter
σ^{min}	Threshold score
$\sigma(r, q', T)$	Total score of a result
q_0	User original query
q'	A candidate query
\mathcal{Q}	A set of candidate queries
\mathcal{S}	A set of Keyword Inverted lists
\mathcal{B}	Candidate query batch
\mathcal{R}	A set of results
\mathcal{R}^*	A set of top-k results
$\Delta(r_1, r_2)$	The score difference between r_1 and r_2
r	A result
v_{slca}	A subtree result root
m	A match node
n	Number of keywords in a query
m^l	Tightest match node
\mathcal{P}	An execution plan
$c(\mathcal{B})$	Cost of an execution plan
T	XML data
$d(r, T)$	Number of edges in a result r
$\theta(r, T)$	Cohesiveness score of a result r
\mathcal{K}	A set of candidate keywords
k	A query keyword

Tightest SLCA Subtree Result: For an SLCA node, there may exist several subtree results. This is because that under an SLCA node v_{slca} , we may find several match nodes $\{m_i^j\}$ for the keyword k_i , ($1 \leq i \leq n$, $1 \leq j \leq n_i$), where n_i is the number of match nodes for k_i under v_{slca} . Let $m_i^{l_i}$ be the closest match node from $\{m_i^j\}$ to v_{slca} for k_i ($1 \leq i \leq n$), then we get the the tightest subtree result $r = (v_{slca}, \{m_1^{l_1}, \dots, m_i^{l_i}, \dots, m_n^{l_n}\})$.

For example, for $SLCA(q) = [0.2]$ in Fig. 2, there are two subtree results, (c) and (d). The tightest subtree result is (c) $r = ([0.2], \{[0.2.1.0.0], [0.2.3.0.0]\})$.

We argue to return only the tightest SLCA subtree results to the user as these results match the user's search intention better than the results containing the sparsely distributed keyword match nodes under v_{slca} . That is, a result is more likely to be meaningful when the result subtree is more tight and cohesive (for survey [14], [12]).

2.2 Problem Statement

Definition 1 (No-Match Problem) Given a keyword query $q_0 = \{k_1, k_2, \dots, k_n\}$ on T , if $\exists k_i \in q_0$ such that $S_i = \emptyset$, we say that query q_0 has a no-match problem over k_i .

If a user submits a keyword query q_0 that has a no-match problem, traditional systems return an empty result set. However, the missing keyword k_i that causes

the no-match problem may have semantic counterparts in the data source T which may produce results the user might be interested in, if the candidate keywords are sufficiently similar to $k_i \in q_0$. We use \mathcal{K}_i to denote the list of candidate keywords that can be used to replace $k_i \in q_0$.

Example 5 . Consider the keyword query $q_0 = \{Jack, lecturer, class\}$ presented in Example 1. It is easy to verify that the keywords $k_2 = lecturer$ and $k_3 = class$ cause a no-match problem for q_0 . Candidate keywords that can be used instead of k_2 and k_3 for q_0 are: $\mathcal{K}_2 = \{academic, full professor\}$ and $\mathcal{K}_3 = \{course, grade, event\}$.

Definition 2 (No-But-Semantic-Match Problem)

Given a keyword query $q_0 = \{k_1, k_2, \dots, k_n\}$ with a no-match problem on T , i.e., $\exists k_i \in q_0$ such that $S_i = \emptyset$, but $\mathcal{K}_i \neq \emptyset$, then we say that q_0 has a no-but-semantic-match problem over k_i .

The no-but-semantic-match problem is a special case of the no-match problem. The problem can be addressed in the following way: (a) find a candidate keyword list \mathcal{K}_i that can be used to replace $k_i \in q_0$; (b) generate candidate queries q' for q_0 by replacing k_i with $k'_i \in \mathcal{K}_i$; (c) execute q' in the data source T to produce the semantically related results \mathcal{R} for q_0 ; (d) score and rank the results $r \in \mathcal{R}$ to return only the top quality results to the user for evaluation.

Example 6 . Consider the keyword query presented in Example 1. The query q_0 has no-but-semantic-match problem over $k_2 = lecturer$ and $k_3 = class$. The semantic counterparts for k_2 and k_3 are: $\mathcal{K}_2 = \{academic, full professor\}$ and $\mathcal{K}_3 = \{course, grade, event\}$. These candidate keywords are combined with the rest of the keywords to generate semantically related candidate queries for q_0 . The generated candidate queries are: $q_1 = \{Jack, academic, course\}$, $q_2 = \{Jack, academic, grade\}$, $q_3 = \{Jack, full professor, course\}$, $q_4 = \{Jack, full professor, grade\}$, $q_5 = \{Jack, academic, event\}$, and $q_6 = \{Jack, full professor, event\}$.

We use \mathcal{Q} to denote the list of candidate queries. As mentioned before, the candidate queries $q' \in \mathcal{Q}$ need to be executed against T to produce the semantically related result set \mathcal{R} for q_0 . We know that these candidate queries q' are generated by replacing k_i with $k'_i \in \mathcal{K}_i$. However, not all candidate keywords $k'_i \in \mathcal{K}_i$ are semantically similar to the user given keyword $k_i \in q_0$ and also, not all semantically related results $r \in \mathcal{R}$ are meaningful to the same degree. Therefore, we need to score the produced results $r \in \mathcal{R}$, denoted by $\sigma(r, q', T)$, as given as follows:

$$\sigma(r, q', T) = sim(q_0, q') \times coh(r, T) \quad (1)$$

where, r is a result for the candidate query q' , the $sim(q_0, q')$ measures the similarity of q' with q_0 and $coh(r, T)$ measures the cohesiveness of r in T . The rank of a result $r \in \mathcal{R}$ is calculated as follows:

$$rank(r, T) = |\{r' | \sigma(r', q'', T) > \sigma(r, q', T)\}| + 1 \quad (2)$$

Definition 3 Top-k Semantically Related Results.

Given a keyword query $q_0 = \{k_1, k_2, \dots, k_n\}$ on T , having the no-but-semantic-match problem, we want to discover k results from \mathcal{R} that maximizes the scoring function given in Eq. 1 or in terms of ranking the results $\{r | rank(r, T) \leq k\}$.

3 Our Approach

We propose a two phase approach to solve the no-but-semantic-match problem in XML data T . The schematic diagram of our approach is illustrated in Fig. 3. In the first phase, the semantic counterparts for the non-mapped keywords of the user query are extracted from the ontological knowledge base. Next, the candidate queries are generated by replacing the non-mapped keywords with their semantic counterparts and the similarities between the candidate queries and the original query are computed. In the second phase, the candidate queries are executed against the data source T . The results are scored based on Eq. 1 and finally, only the top-k results are presented to the user for evaluation. The results are scored based on the followings: (a) similarity of the candidate queries to the user given query; and (b) the cohesiveness of the results.

As the candidate queries are generated using the ontological knowledge base, we use the *ontological similarity* of the candidate query to the user given query as the measure of similarity for the first parameter. The details for computing this similarity is presented in section 3.1. The details for computing the *cohesiveness* of the results is presented in section 3.2. Since there are a number of candidate queries that should be executed against the data source T and each candidate query may have several results that needs to be scored, we propose efficient pruning techniques to avoid unnecessary computations and terminate early. The pruning ideas and the details of our candidate query processing technique are presented in Section 3.3. We also propose a *batch* query processing technique to speed up the computations further by sharing the computations among the candidate queries, which is described in Section 4.

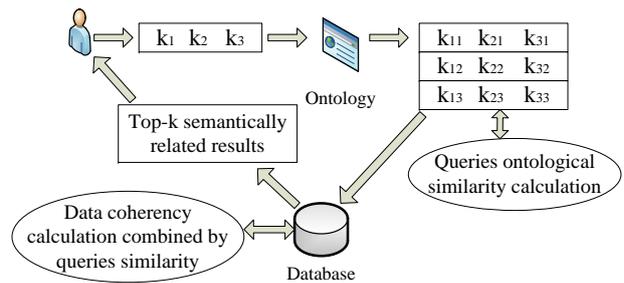


Fig. 3: Schematic diagram of our approach for solving no-but-semantic-match problem in XML data

3.1 Candidate Queries

This section describes how to generate the candidate queries \mathcal{Q} for the user query q_0 and compute their similarity to q_0 .

3.1.1 Generating Candidate Queries

To adhere closely to the user's intentions, the candidate keywords $k'_i \in \mathcal{K}_i$ must be as close as possible to the non-mapped keywords $k_i \in q_0$. In this study we use WordNet, which is widely used in the literature [5] for finding semantic counterparts for $k_i \in q_0$. We categorize our semantic candidate keywords derived from WordNet into four groups [18]: (a) synonyms denoted as $Syn(k_i)$, (b) coordinate terms denoted as $Cot(k_i)$, (c) hyponyms denoted as $Hpo(k_i)$, and (d) hypernyms denoted as $Hpe(k_i)$. The candidate keyword list \mathcal{K}_i for a keyword $k_i \in q_0$ contains all types of ontological counterparts as shown in Eq. 3.

$$\mathcal{K}_i = Syn(k_i) \cup Cot(k_i) \cup Hpo(k_i) \cup Hpe(k_i) \quad (3)$$

However, not all candidate keywords extracted from the ontological knowledge base are available in T . The keyword list has to be reduced to the candidates which have direct mapping in the data source. To do this, an inverted keyword list using hash indices can be queried in $\mathcal{O}(1)$ time. Finally, the candidate queries \mathcal{Q} are generated by replacing the non-mapped keywords $k_i \in q_0$ with each of their semantic counterparts $k'_i \in \mathcal{K}_i$ in turn.

3.1.2 Measuring Candidate Query Similarity

In order to measure the similarity between a candidate query $q' \in \mathcal{Q}$ and the user's original query q_0 , firstly we measure the individual similarity between the candidate keyword $k'_i \in q'$ and the corresponding non-mapped keyword $k_i \in q_0$ using Wu and Palmer's metric

[25]. This metric establishes the depths of both keywords and their least common subsumer (LCS) according to the WordNet structure and produces the degree of similarity between these two keywords, $SimWP(k_i, k'_i)$, as shown in Eq. 4.

$$SimWP(k_i, k'_i) = \frac{2 \times dep(LCS)}{dep(k_i) + dep(k'_i)} \quad (4)$$

where $dep(k_i)$ returns the depth of the keyword k_i in the WordNet structure. This metric is symmetric, i.e., $SimWP(k_i, k'_i) = SimWP(k'_i, k_i)$. However, WordNet has a hierarchical structure. That is, the candidate keyword $k'_i \in q'$ could be a more special type (e.g., hyponyms) or a more general type (e.g., hypernyms) for the non-mapped keyword $k_i \in q_0$ in WordNet. Therefore, we incorporate the specialization/generalization aspect of $k_i \in q_0$ into the Wu and Palmer similarity metric as given as follows:

$$DSim(k_i, k'_i) = \frac{dep(k'_i)}{max(dep(k_i, k'_i))} \times SimWP(k_i, k'_i) \quad (5)$$

where $DSim(k_i, k'_i)$ is the directional similarity of keyword $k_i \in q_0$ to keyword $k'_i \in q'$. The directional similarity $DSim(k_i, k'_i)$ penalizes the more general keyword types of $k_i \in q_0$ by weighting $SimWP(k_i, k'_i)$ with $\frac{dep(k'_i)}{max(dep(k_i, k'_i))}$. Finally, the similarity of q' to the original query q_0 is computed by considering all the replacements in q' as follows:

$$\lambda(q_0, q') = \prod_{i=1}^n DSim(k_i \in q_0, k'_i \in q') \quad (6)$$

where n is the number of keywords in q_0 that have been replaced to generate q' ($1 \leq n \leq |q|$).

Example 7. Consider the keyword query $q_0 = \{Jack, lecturer, class\}$ presented in Example 1. Here, the second and the third keywords cause the no-match problem for q_0 . The candidate keyword list for these two non-mapped keywords are: $\mathcal{K}_2 = \{academic : 0.91, full\ professor : 0.84\}$ and $\mathcal{K}_3 = \{course : 1, grade : 1, event : 0.35\}$, where each candidate keyword is labeled with their corresponding $DSim$ scores. Now, the candidate queries are generated by replacing the non-mapped keywords in q_0 with their candidate keywords and scored as follows:

$q_1 = \{Jack, academic, course\}, \lambda(q_0, q_1) = 0.91 \times 1 = 0.91$,
 $q_2 = \{Jack, academic, grade\}, \lambda(q_0, q_2) = 0.91 \times 1 = 0.91$,
 $q_3 = \{Jack, fullprof, course\}, \lambda(q_0, q_3) = 0.84 \times 1 = 0.84$,
 $q_4 = \{Jack, fullprof, grade\}, \lambda(q_0, q_4) = 0.84 \times 1 = 0.84$,
 $q_5 = \{Jack, academic, event\}, \lambda(q_0, q_5) = 0.91 \times 0.35 = 0.31$, and
 $q_6 = \{Jack, fullprof, event\}, \lambda(q_0, q_6) = 0.84 \times 0.35 = 0.29$.

From the above, it is easy to verify that q_1 and q_2 are the most similar candidate queries to the original query.

3.2 Cohesiveness of Results

There can be potentially many candidate queries with a large number of results and not all results $r \in \mathcal{R}$ are meaningful to the same degree. In XML, if the match nodes in the result r are near to each other, the result is considered to be more cohesive. Intuitively, when a result subtree is more cohesive, it is more likely to be relevant and meaningful (for survey [14], [12]). To measure the cohesiveness of a result r , we firstly compute the distance between each match node m and the root v_{slca} of the result subtree r . Then, we compute the overall distance of a result r w.r.t. the data source T as given as follows [12]:

$$d(r, T) = \sum (l_m - l_{v_{slca}}) \quad (7)$$

where r is the result subtree, v_{slca} denotes the root of the result r , m is the match node, l_m is the level of m , and $l_{v_{slca}}$ is the level of the root in r . Clearly, the larger this distance is, the lower the cohesiveness score for the result r should be. Therefore, we compute the cohesiveness of a result subtree r w.r.t. the data source T as given as follows [12]:

$$\theta(r, T) = \frac{1}{\log_{\alpha}(d(r, T) + 1) + 1} \quad (8)$$

where α is the tuning parameter by which the user can trade off between the similarity of the candidate queries and the cohesiveness of the results. If we set α to a larger value, the sensitivity to the cohesiveness of the results gets smaller. That is, the total score σ of a result r is more dependent on the similarity of the query $\lambda(q_0, q')$ than its cohesiveness.

Example 8. Consider the candidate keyword queries \mathcal{Q} in Example 7. The result subtrees of these queries are illustrated in Fig. 4. Using $\alpha = 4$, we compute the cohesiveness of the results as follows:

$d(r_1, T) = 7, \theta(r_1, T) = \frac{1}{2.5} = 0.4$,
 $d(r_2, T) = 10, \theta(r_2, T) = \frac{1}{2.72} = 0.36$,
 $d(r_3, T) = 8, \theta(r_3, T) = \frac{1}{2.58} = 0.38$,
and $d(r_4, T) = 11, \theta(r_4, T) = \frac{1}{2.79} = 0.35$.

3.3 Processing of Candidate Queries

A naïve approach to processing the no-but-semantic-match query q_0 first generates the candidate queries

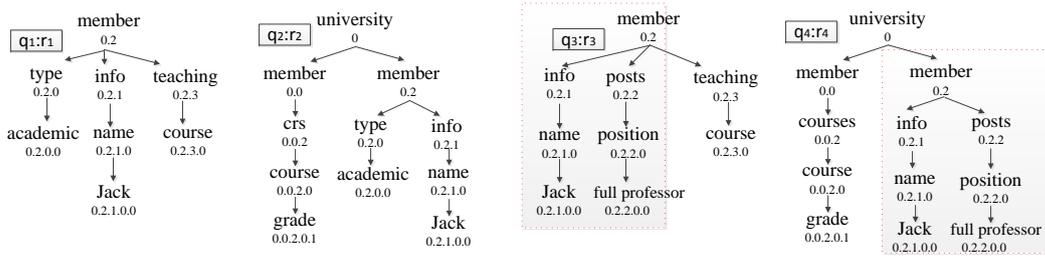


Fig. 4: Subtree results for candidate queries \mathcal{Q} in Example 7 after executing them against data given in Figure 1.

\mathcal{Q} and then computes all semantically related results $r \in \mathcal{R}$ by executing the queries $q' \in \mathcal{Q}$ against the data source T . Then it applies Eq. 6 to the results \mathcal{R} to establish the similarity of the corresponding candidate query q' to the original query q_0 and determines the cohesiveness in the data source T according to Eq. 8. The results are then sorted based on their total score σ (Eq. 1) to obtain the top- k ranked (Eq. 2) semantically related results.

Assume that the initial query $q_0 = \{k_1, \dots, k_n\}$ has the no-but-semantic-match problem for all $k_i \in q_0$, $|\mathcal{Q}|$ is the maximal number of candidate queries produced for q_0 , d is the depth of the tree T , $|S|$ and $|S_1|$ are the maximal and minimal sizes of the inverted keyword lists for the semantic counterparts k'_i , respectively and $|\mathcal{R}|$ is the maximal number of semantically related results for the candidate queries in \mathcal{Q} , then the complexity of the naive approach becomes $|\mathcal{Q}| \times nd|S_1| \log |S| + |\mathcal{R}| \log |\mathcal{R}|$. However, both $|\mathcal{Q}|$ and $|\mathcal{R}|$ could be potentially large, which makes the naive approach impractical. We propose two efficient pruning techniques, called the *inter-query* pruning and *intra-query* pruning to significantly reduce the sizes of \mathcal{Q} and \mathcal{R} , respectively.

3.3.1 Inter-Query Pruning

Assume the candidate queries $\mathcal{Q} = \{q_1, q_2, \dots, q_l, q_{l+1}, \dots, q_{|\mathcal{Q}|}\}$ are sorted based on their similarities $\lambda(q_0, q_i)$ to the original query q_0 . We obtain the following lemma.

Lemma 1 Assume R^* is the k results of $\mathcal{Q}' = \{q_1, q_2, \dots, q_l\}$ and σ^{min} is the min-score of the results R^* . Then, we can stop processing the rest of the candidate queries $\mathcal{Q}'' = \{q_{l+1}, \dots, q_{|\mathcal{Q}|}\}$ if $\lambda(q_0, q_{l+1}) < \sigma^{min}$.

Proof Assume that r is a min-scored result in R^* for $q' \in \mathcal{Q}'$, i.e., $\sigma^{min} = \sigma(r, q', T)$ and r' is a result of the candidate query q_{l+1} whose similarity score is higher than any result of the queries $\mathcal{Q}'' = \{q_{l+1}, \dots, q_{|\mathcal{Q}|}\}$. Now, assume that $\sigma(r', q_{l+1}, T) > \sigma(r, q', T)$. We prove that this can not happen if $\lambda(q_0, q_{l+1}) < \sigma^{min}$. To be scored higher than r , r' must satisfy the following: $\lambda(q_0, q_{l+1}) > \frac{\sigma^{min}}{\theta(r', T)}$. However, the highest possible value of θ for any

result in T is 1. By putting this into the above, we get $\lambda(q_0, q_{l+1}) > \sigma^{min}$, which contradicts the assumption. Therefore, R^* consists of the top- k semantically related results according to Def. 3 whose ranks are $\leq k$.

Example 9. Consider the candidate queries \mathcal{Q} given in Example 7. If we want to present the top-1 result to the user, and after processing the candidate queries up to q_4 we get $\sigma^{min} = 0.91 \times 0.4 = 0.36$, then we do not need to process q_5 as $\lambda(q_0, q_5) = 0.31 < \sigma^{min}$. Thus, from q_5 to the end of the list of \mathcal{Q} , no queries can score higher than σ^{min} and therefore, we can stop processing them.

3.3.2 Intra-Query Pruning

Although the *inter-query pruning* technique does not execute all of the candidate queries $q' \in \mathcal{Q}$ against T , it employs the pruning technique only in the first phase of the framework. That is, once we start processing a candidate query q' , we compute all of its results. Consider the candidate queries \mathcal{Q} of q_0 given in Example 7 and assume that the user requests only the top-1 result for q_0 . Also, assume that the candidate queries in \mathcal{Q} are sorted based on their similarities with q_0 and we have already processed the candidate queries from q_1 to q_5 . The current top-1 result is r_1 (see in Fig. 4) and σ^{min} is 0.36. Now, while processing the candidate query q_3 , we can discard the result r_3 of q_3 (as shown in Fig. 4) while generating it. That is, while reading through the keyword inverted lists S_{Jack} , $S_{fullprofessor}$ and S_{course} for q_3 , we can partially compute r_3 consisting of the keywords $\{Jack, full\ professor\}$ for q_3 only (as highlighted in Fig. 4), denoted by r_3^p , and compare the score 0.34 of r_3^p with the current σ^{min} , we can decide that the complete r_3 consisting of keywords $\{Jack, full\ professor, course\}$, denoted by r_3^c , can never outrank r_1 as $\theta(r_3^c, T) \leq \theta(r_3^p, T)$. The same occurs while computing r_4 of q_4 and we can discard r_4 before generating the ultimate result. We call the above query pruning technique as the *intra-query pruning*.

Algorithm 1: The Framework

Input : User Query q_0 , Tuning Parameter α ,
Keyword Inverted Lists \mathcal{S}

Output: Top-k Semantically Related Results \mathcal{R}^*

```

1  $\mathcal{Q} \leftarrow \text{generateCandidateQueries}(q_0)$ ;
2 while  $q' \leftarrow \mathcal{Q}.\text{getNext}() \neq \text{null}$  do
3    $q'.\text{sim} \leftarrow \lambda(q_0, q')$ ; // according to Eq.6
4    $\mathcal{Q} \leftarrow \text{sortCandidates}(\mathcal{Q})$ ; // based on  $q'.\text{sim}$ 
5    $\mathcal{R}^* \leftarrow \text{null}$ ; //  $\mathcal{R}^*$  is a min heap
6    $\sigma^{\text{min}} \leftarrow \text{MAXVAL}$ ; // max value
7 while  $q' \leftarrow \mathcal{Q}.\text{getNext}() \neq \text{null}$  do
8   if  $|\mathcal{R}^*| = k$  and  $q'.\text{sim} < \sigma^{\text{min}}$  then
9     break; // inter-query pruning
10  if  $|\mathcal{R}^*| = k$  then
11     $\text{root} \leftarrow \mathcal{R}^*.\text{root}()$ ;  $\sigma^{\text{min}} \leftarrow \text{root}.\text{score}$ ;
12   $\mathcal{S}' \leftarrow \{\}$ ;
13  foreach  $k_i \in q'$  do
14     $S_i \leftarrow \text{retrieveKeywordInvertedList}(k_i, \mathcal{S})$ ;
15     $\mathcal{S}' \leftarrow \mathcal{S}' \cup S_i$ ;
16   $\mathcal{R}^* \leftarrow \text{processQuery}(\mathcal{R}^*, \sigma^{\text{min}}, \alpha, q', \mathcal{S}')$ ;
17 return  $\mathcal{R}^*$ 

```

3.3.3 The Framework

Algorithm 1 presents the framework for processing the no-but-semantic-match query q_0 submitted by the user. First, it generates the candidate queries \mathcal{Q} for q_0 as explained in Section 3.1.1, shown on line 1. The lines 2-4 compute the similarity between the candidate queries \mathcal{Q} and the user query q_0 as explained in Section 3.1.2 and sort them. Then, a min-heap is initialized with \mathcal{R}^* to *null* and the min-score σ^{min} to *MAXVAL* in lines 5-6. In lines 8-9, we stop processing the candidate queries in \mathcal{Q} as soon as we find a query $q' \in \mathcal{Q}$ if $|\mathcal{R}^*| = k$ and $q'.\text{sim} < \sigma^{\text{min}}$. Otherwise, if $|\mathcal{R}^*| = k$, we update σ^{min} by reading the root entry of the heap \mathcal{R}^* and adding its score $\text{root}.\text{score}$ to σ^{min} in lines 10-11. Then, for each keyword $k_i \in q'$ we retrieve their corresponding inverted lists and pass it to the *processQuery* method (which is explained in detail in the following section) in line 16 to retrieve the results of q' and insert the eligible results into \mathcal{R}^* .

3.3.4 The processQuery Method

In order to process the no-but-semantic-match query q_0 , we need to execute each candidate keyword query $q' \in \mathcal{Q}$ against the data source T in the *processQuery* method. There are two benchmark algorithms in the literature to compute the keyword query results on XML data T as given as follows: (a) scan eager [26] and (b) anchor based [23] algorithms. However, these two benchmark algorithms are not readily available to implement our *processQuery* method. These algorithms only find the root of the subtree results in T , but ignore the distribution of the keyword match nodes in the

Algorithm 2: SE-QP

Input : $\mathcal{R}^*, \sigma^{\text{min}}, \alpha, q', \mathcal{S}'$

Output: Result Set: \mathcal{R}^*

```

1  $\text{sortLists}(\mathcal{S}')$ ;  $r \leftarrow \text{null}$ ;  $v_{\text{slca}} \leftarrow \text{null}$ ;
2 while  $m_1 \leftarrow \text{getNext}(S_1) \neq \text{null}$  do
3    $m_i \leftarrow \text{closest}(m_1, S_i), \forall i \in [2, n]$ ;
4    $v_{\text{slca}}^u \leftarrow \text{lca}(m_1, \dots, m_n)$ ;
5   if  $r \neq \text{null}$  and  $v_{\text{slca}} \not\prec_a v_{\text{slca}}^u$  then
6     for  $i = 1 \rightarrow n$  do
7        $m_i^{l_i} \leftarrow \text{getTight}(S_i, r.\text{cursor}_i)$ ;
8        $d_i \leftarrow \text{getDist}(v_{\text{slca}}, m_i^{l_i})$ ;  $d \leftarrow d + d_i$ ;
9        $r.\text{score} \leftarrow q'.\text{sim} \times \frac{1}{\log_a(d+1)+1}$ ;
10      if  $r.\text{score} < \sigma^{\text{min}}$  then
11        stop reading lists and jump to line 14.
12       $r \leftarrow (v_{\text{slca}}, \{m_i^{l_i}, \forall i \in [1, n]\})$ ;
13      update top-k list  $\mathcal{R}^*$  with  $r$ ;
14       $r \leftarrow \text{null}$ ;
15      if  $v_{\text{slca}}^u \not\prec_a v_{\text{slca}}$  then
16         $v_{\text{slca}} \leftarrow v_{\text{slca}}^u$ ;  $r.\text{add}(S_i.\text{cursor}, \forall i \in [1, n])$ ;
17      if  $v_{\text{slca}}^u \not\prec_a v_{\text{slca}}$  then
18         $r \leftarrow (v_{\text{slca}}^u, \{m_i^{l_i}, \forall i \in [1, n]\})$ ;
19        score  $r$  and update  $\mathcal{R}^*$  with  $r$  if  $r.\text{score} > \sigma^{\text{min}}$ ;
20 return  $\mathcal{R}^*$ 

```

subtree. To implement our *processQuery* method with these benchmark algorithms, we need to address the following issues which are specific to our problem: (a) finding the tightest nodes under the confirmed SLCA root v_{slca} and (b) scoring the result partially based on its candidate query similarity and cohesiveness to apply *intra-query* pruning. We propose two techniques to implement the *processQuery* method based on the benchmark algorithms as follows:

1. Scan Eager based Query Processing (SE-QP) and
2. ANchor based Query Processing (AN-QP).

SE-QP Algorithm. Like scan-eager algorithm[26], SE-QP firstly sorts the inverted lists of the keywords in q' . Then, it picks a match node m_1 from the shortest inverted list S_1 and then, finds the closest match nodes to m_1 from other lists to compute the result root v_{slca} . However, the tightest subtree result computation and result scoring is delayed until we can confirm that this v_{slca} can be an actual SLCA node. Therefore, the cursor of each of the inverted list is retained until we decide that this v_{slca} cannot be an ancestor of any other result roots. Once we confirm that this v_{slca} is an actual SLCA node, we compute the tightest subtree result for it and score the result. While scoring the result, we also apply *intra-query* pruning here. Then, we advance the cursors of all of the lists and continue the above steps until we access all nodes in the list S_1 .

The SE-QP query processing technique is pseudocoded in Algorithm 2. In line 1, we sort the inverted list of all keywords in q' and do the initialization. In lines 2-3,

it finds the match node m_1 from the shortest list S_1 and the match nodes m_i from other lists. In line 4, we find the potential root result v_{slca}^u which then should be confirmed as an actual SLCA node. In line 5, we check v_{slca}^u with the previous result root v_{slca} . If v_{slca} is not an ancestor for v_{slca}^u , denoted as $v_{slca} \not\prec_a v_{slca}^u$, v_{slca} is confirmed as the SLCA result root. Then, the corresponding tightest subtree result for v_{slca} is retrieved. To do so, we scan each list S_i by moving its cursor $r.cursor_i$ backward and forward to find the closest match nodes under v_{slca} , which is implemented in function *getTight* of line 7. For each closest match node m_i^{li} , the distance of m_i^{li} with v_{slca} is computed by *getDist* function in line 8 and the score of the result $r.score$ is computed partially in lines 8-9. We stop scanning other lists if the partial score cannot beat σ^{min} (intra-query pruning) and jump to line 14, which is given in lines 10-11. Otherwise, we keep scanning all lists to compute the tightest subtree result and the ultimate score of r for the v_{slca} . We update the min heap \mathcal{R}^* by this result r in line 13. Now, we update v_{slca} with the current result root v_{slca}^u if $v_{slca}^u \not\prec_a v_{slca}$ in lines 15-16. In lines 17-19, if the last result root node is an actual SLCA, the similar steps are conducted to score its tightest subtree result and if promising, is used to update \mathcal{R}^* .

AN-QP Algorithm. Like scan-eager algorithm[26], SE-QP performs worse when the inverted lists have similar sizes (e.g., match node distribution). Also, it incurs many redundant computations when the data distribution is skewed in T . For example, if the match nodes are mostly distributed in one part of the XML tree in an inverted list, SE-QP reads all the nodes in S_1 and computes their LCAs to finalize the corresponding SLCA nodes. However, lots of these nodes can be skipped because they are far from the nodes in other inverted lists and cannot create SLCA nodes.

In order to skip the non-promising match nodes, like [23], AN-QP considers only the *anchor* match nodes for computing SLCA nodes. A set of match nodes $M = \{m_1, \dots, m_n\}$ for q' is said to be *anchored* by a match node $m_a \in M$ if for each $m_i \in M \setminus \{m_a\}$, $m_i = \text{closest}(m_a, S_i)$, where *closest*(m_a, S_i) returns the match nodes in the list S_i which is closest to the node m_a [23]. Unlike SE-QP, the *anchor* match node m_a is picked from among inverted lists (not necessarily from the shortest one) so that it can maximize the skipping of redundant computations. Similar to SE-QP, AN-QP first extracts the SLCA result root v_{slca} and thereafter, finds the tightest subtree result under v_{slca} and partially score the result to apply *intra-query* pruning.

The AN-QP technique is pseudocoded in Algorithm 3. Line 2 finds the anchor node m_a from the the inverted lists $\{S_1, \dots, S_n\}$, which is implemented in func-

Algorithm 3: AN-QP

Input : $\mathcal{R}^*, \sigma^{min}, \alpha, q', S'$
Output: Result Set: \mathcal{R}^*

```

1  $r \leftarrow null; v_{slca} \leftarrow null;$ 
2  $m_a \leftarrow \text{getAnchor}(\{\text{getNext}(S_i), \forall i \in [1, n]\});$ 
3 while  $m_a \neq null$  do
4    $m_i \leftarrow \text{closest}(m_a, S_i), \forall i \in [1, n] \& i \neq a;$ 
5    $v_{slca}^u \leftarrow \text{lca}(m_1, \dots, m_n);$ 
6   if  $r \neq null$  and  $v_{slca} \not\prec_a v_{slca}^u$  then
7     for  $i = 1 \rightarrow n$  do
8        $m_i^{li} \leftarrow \text{getTight}(S_i, r.cursor_i);$ 
9        $d_i \leftarrow \text{getDist}(v_{slca}, m_i^{li}); d \leftarrow d + d_i$ 
10       $r \cdot \text{score} \leftarrow q' \cdot \text{sim} \times \frac{1}{\log_a(d+1)+1};$ 
11      if  $r \cdot \text{score} < \sigma^{min}$  then
12        | stop reading lists and jump to line 15.
13       $r \leftarrow (v_{slca}, \{m_i^{li}, \forall i \in [1, n]\});$ 
14      update top-k list  $\mathcal{R}^*$  with  $r;$ 
15     $r \leftarrow null;$ 
16    if  $v_{slca} \not\prec_a v_{slca}^u$  then
17      |  $v_{slca} \leftarrow v_{slca}^u; r.add(S_i.cursor), \forall i \in [1, n];$ 
18       $m_a \leftarrow \text{getAnchor}(\{\text{getNext}(S_i), \forall i \in [1, n]\});$ 
19  if  $v_{slca}^u \not\prec_a v_{slca}$  then
20    |  $r \leftarrow (v_{slca}^u, \{m_i^{li}, \forall i \in [1, n]\});$ 
21    score  $r$  and update  $\mathcal{R}^*$  with  $r$  if  $r \cdot \text{score} > \sigma^{min};$ 
22 return  $\mathcal{R}^*$ 

```

tion *getAnchor*. The potential result root v_{slca}^u is computed after finding the closest nodes $m_i \in S_i$ to m_a , $\forall i \in [1, n]$ and $i \neq a$ in lines 4-5. Here, if $v_{slca} \not\prec_a v_{slca}^u$, v_{slca} is confirmed as the SLCA result root as given in line 6. Therefore, we retrieve the closest match nodes $m_i^{li}, \forall i \in [1, n]$ under v_{slca} . Similar to SE-QP, the function *getTight* of line 8, scans each list S_i by moving its cursor $r.cursor_i$ backward and forward to compute the tightest subtree result for v_{slca} . Then, we find the distance of each m_i^{li} with v_{slca} by the function *getDist* and add it to the total distance of r in line 9. Then, we compute the score of the result $r.score$ partially in line 10. Here, we apply intra-query pruning if the partial score cannot beat σ^{min} and jump to line 15, which is given in lines 11-12. Otherwise, we compute the ultimate score of the result by scanning all lists to retrieve the tightest subtree result under v_{slca} and update r with its corresponding data in line 13. In line 14, the promising result r is inserted into \mathcal{R}^* . The current result root v_{slca}^u is saved into v_{slca} , if $v_{slca}^u \not\prec_a v_{slca}$ in lines 16-17. We update the anchor node m_a in line 18 for the next iteration. We check the last result root node and retrieve its tightest subtree result if it is an actual SLCA node in lines 19-20. Finally, we score it and update \mathcal{R}^* with it if $r \cdot \text{score} > \sigma^{min}$ in line 21.

B ₁	academic	course	Jack	q ₁
B ₂	Jack	academic	grade	q ₂
			event	q ₅
B ₃	Jack	full professor	course	q ₃
			grade	q ₄
			event	q ₆

(a) Plan \mathcal{P}_1 , $c(\mathcal{P}_1) = 38$

B ₁	Jack	course	academic	q ₁
			full professor	q ₃
B ₂	Jack	grade	academic	q ₂
			full professor	q ₄
B ₃	Jack	event	academic	q ₅
			full professor	q ₆

(b) Plan \mathcal{P}_2 , $c(\mathcal{P}_2) = 39$

B ₁	Jack	academic	course	q ₁
			grade	q ₂
			event	q ₅
B ₂	Jack	full professor	course	q ₃
			grade	q ₄
			event	q ₆

(c) Plan \mathcal{P}_3 , $c(\mathcal{P}_3) = 34$

Fig. 5: A part of the possible execution plans for processing the candidate queries in \mathcal{Q} .

4 Batch Processing

This section investigates a more efficient method for processing the candidate queries \mathcal{Q} obtained from the initial query q_0 with the no-but-semantic-match problem. As the candidate queries in \mathcal{Q} are generated by replacing the keywords of the initial query, they usually share a subset of keywords. It is possible to compute the results of these shared subsets of keywords among the queries in \mathcal{Q} and then merge the results of the shared part with the exclusive part of each query $q' \in \mathcal{Q}$ [28]. However, two challenges arise here: (a) finding the groups of queries, called *batches*, which share a subset of keywords among them and can be executed efficiently; and (b) finding the tightest SLCA results for the shared part which can be ultimately merged with the exclusive part of each candidate query in a batch.

4.1 Constructing the Candidate Query Batch

Given two candidate queries $q_1, q_2 \in \mathcal{Q}$, assume that $\mathcal{K}^s(q_1, q_2)$ denotes the set of keywords shared by them, i.e., $\mathcal{K}^s(q_1, q_2) \subseteq q_1, q_2$. We can achieve the best performance by putting q_1 and q_2 in a batch if they share the maximal set of keywords, i.e., $|\mathcal{K}^s(q_1, q_2)| = |q_1| - 1 = |q_2| - 1$ [28]. A candidate query batch is defined as given below:

Definition 4 A candidate query batch, denoted by \mathcal{B} , is a subset of \mathcal{Q} such that the following conditions hold: (a) $\forall q_1, q_2 \in \mathcal{B}$, $|\mathcal{K}^s(q_1, q_2)| = |q_1| - 1 = |q_2| - 1$; (b) $\forall q_1, q_2, q_3 \in \mathcal{B}$, $\mathcal{K}^s(q_1, q_2) = \mathcal{K}^s(q_2, q_3) = \mathcal{K}^s(q_1, q_3)$; and (c) $1 \leq |\mathcal{B}| \leq |\mathcal{Q}|$.

To execute the candidate queries in \mathcal{Q} , we have to construct the set of batches that can cover all queries in \mathcal{Q} . We call the set of candidate query batches that cover the queries in \mathcal{Q} an *execution plan*, which is defined below:

Definition 5 An execution plan, denoted by \mathcal{P} , is a set of candidate query batches such that: (a) $\mathcal{Q} = \bigcup_{i=1}^{|\mathcal{P}|} \mathcal{B}_i \in \mathcal{P}$ and (b) $\forall \mathcal{B}_1, \mathcal{B}_2 \in \mathcal{P}$, $\mathcal{B}_1 \cap \mathcal{B}_2 = \emptyset$.

An execution plan \mathcal{P} has its evaluation cost which is the summation of the execution costs of its constituent batches. The execution cost of a batch \mathcal{B} directly depends on the inverted keyword lists which have to be accessed. Assume that \mathcal{K}^u is the set of keywords of all candidate queries in a batch \mathcal{B} that has not been covered by \mathcal{K}^s , i.e., $\bigcup_{q_1 \in \mathcal{B}} q_1 \setminus \mathcal{K}^s$. We estimate the cost of executing \mathcal{B} , denoted by $c(\mathcal{B})$ as:

$$c(\mathcal{B}) = \sum_{k_1 \in \mathcal{K}^s} |S_{k_1}| + \sum_{k_2 \in \mathcal{K}^u} (|\min(S_{\mathcal{K}^s})| + |S_{k_2}|) \quad (9)$$

where $\min(S_{\mathcal{K}^s})$ returns the shortest inverted keyword list size among the keywords in \mathcal{K}^s . However, there exist many plans for \mathcal{Q} as shown in Fig. 5. The optimal plan has the least cost. Discovering this optimal plan is a combinatorial optimization problem as there are many ways of constructing the candidate query batches from \mathcal{Q} . Here, we propose a *greedy approach* for discovering a sub-optimal plan which consists of the following steps: (a) retrieve the topmost similar query $q_1 \in \mathcal{Q}$ to q_0 ; (b) construct all plausible batches for q_1 as follows: (i) remove a keyword $k_1 \in q_1$ and construct \mathcal{K}^s as $q_1 \setminus k_1$; (ii) retrieve all $q_2 \in \mathcal{Q}$ such that $\mathcal{K}^s \subset q_2$; (iii) insert q_1 and all q_2 into a plausible batch \mathcal{B}_1 ; (c) make the batch \mathcal{B}_1 as the actual batch that has the least unit cost $\frac{c(\mathcal{B}_1)}{|\mathcal{B}_1|}$; (d) remove all queries \mathcal{B}_1 from \mathcal{Q} ; and (e) repeat the above steps until \mathcal{Q} is empty.

Example 10. Consider the candidate queries \mathcal{Q} presented in Example 7 and the sizes of the inverted lists as follows: $|S_{Jack}| = 3$, $|S_{academic}| = 1$, $|S_{fullprofessor}| = 1$, $|S_{course}| = 6$, $|S_{grade}| = 2$, $|S_{event}| = 2$.

We start with the topmost query $q_1 \in \mathcal{Q}$ and remove one keyword at a time from q_1 to create the plausible candidate query batches as follows:

$\mathcal{B}_{1.1}$ ($\mathcal{K}^s = \{Jack, academic\}$, $\mathcal{K}^u = \{course, grade, event\}$)
 $\mathcal{B}_{1.2}$ ($\mathcal{K}^s = \{Jack, course\}$, $\mathcal{K}^u = \{academic, fullprofessor\}$)
 $\mathcal{B}_{1.3}$ ($\mathcal{K}^s = \{academic, course\}$, $\mathcal{K}^u = \{Jack\}$)

We estimate the costs of the candidate query batches as follows: $\frac{c(\mathcal{B}_{1.1})}{|\mathcal{B}_{1.1}|} = 5.6$, $\frac{c(\mathcal{B}_{1.2})}{|\mathcal{B}_{1.2}|} = 8.5$, $\frac{c(\mathcal{B}_{1.3})}{|\mathcal{B}_{1.3}|} = 11$

Therefore, $\mathcal{B}_{1.1}$ is the initial candidate query batch. After excluding the queries in $\mathcal{B}_{1.1}$ from \mathcal{Q} , the next topmost similar candidate query q_3 is considered and the

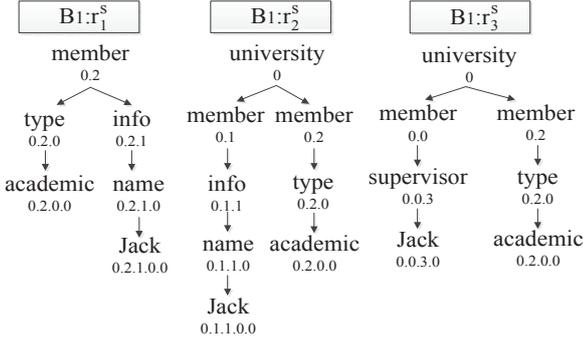


Fig. 6: Shared part processing for query batch \mathcal{B}_1 of \mathcal{P}_3 .

following plausible batches are constructed:

$\mathcal{B}_{2.1}(\mathcal{K}^s = \{Jack, fullprofessor\}, \mathcal{K}^u = \{course, grade, event\}, \dots, q_{|\mathcal{B}_i|})$

$\mathcal{B}_{2.2}(\mathcal{K}^s = \{Jack, course\}, \mathcal{K}^u = \{fullprofessor\})$

$\mathcal{B}_{2.3}(\mathcal{K}^s = \{fullprofessor, course\}, \mathcal{K}^u = \{Jack\})$

The costs of the above plausible batches are as follows:

$\frac{c(\mathcal{B}_{2.1})}{|\mathcal{B}_{2.1}|} = 5.6$, $\frac{c(\mathcal{B}_{2.2})}{|\mathcal{B}_{2.2}|} = 13$, $\frac{c(\mathcal{B}_{2.3})}{|\mathcal{B}_{2.3}|} = 11$

Hence, $\mathcal{B}_{2.1}$ is the next candidate query batch. Now, if we exclude all queries in $\mathcal{B}_{2.1}$ from \mathcal{Q} , \mathcal{Q} becomes empty and the process stops. Finally, plan \mathcal{P}_3 is our execution plan as shown in Fig. 5(c), which is sub-optimal.

4.2 Processing the Candidate Query Batch

To process a candidate query batch \mathcal{B} , we need to compute the results of the shared part \mathcal{K}^s first. Then, we need to merge these shared part results with the non-shared keywords \mathcal{K}^u for $q'' \in \mathcal{B}$. However, computing the shared part results that can be merged with the unshared part is a non-trivial problem. This is because, the SLCA result roots of the shared part \mathcal{K}^s do not guarantee to be the SLCA result roots for $q'' \in \mathcal{B}$. The result roots of q'' may ascend to higher levels in the tree T when the shared part results are merged with the unshared part $q'' \setminus \mathcal{K}^s$.

Consider the potential shared part results of $\{Jack, academic\}$ for the batch \mathcal{B}_1 of plan \mathcal{P}_3 as presented in Fig. 6. Now, when we merge these results with the keyword *course* for $q_1 \in \mathcal{B}_1$, r_1^s contributes to the final SLCA result root, which is $v_{slca}^{r_1^s} = [0.2]$ (see in Fig. 4). However, when we merge these potential shared part results with the keyword *grade* for $q_2 \in \mathcal{B}_1$, r_3^s contributes to the final SLCA result and the result root ascends to $v_{slca}^{r_3^s} = [0]$ (see in Fig. 4). This indicates that we need to retain r_1^s as well as r_3^s as the actual shared part results, though the root of r_3^s is not the SLCA result root of the shared part $\{Jack, academic\}$, but the root of r_1^s is. Here, we do not need to retain r_2^s as $d(r_2^s, T) = 7 > d(r_3^s, T) = 6$ and both r_2^s and r_3^s share the same root.

Lemma 2 Assume $q'' \in \mathcal{B}$ and v_{slca}^s is the SLCA result root of the shared part of \mathcal{B} . Then, the tightest match nodes of the final result of q'' are under v_{slca}^s or under one of the ancestors of v_{slca}^s .

Therefore, to process a candidate query batch shared part \mathcal{K}^s , we find the closest match nodes under the shared part result root v_{slca}^s as well as under all of its ancestors.

4.2.1 Inter and Intra-Batch Pruning

Assume the candidate queries in a batch \mathcal{B}_i are sorted based on their similarities to q_0 as follows: $\{q_1, q_2, \dots, q_{l-1}, \dots, q_{|\mathcal{B}_i|}\}$. Also, $\mathcal{R}_{\mathcal{B}_i}^s$ is the set of shared part results for \mathcal{B}_i . Then, the lower bound of the overall distances of the results $\mathcal{R}_{\mathcal{B}_i}$ of \mathcal{B}_i , denoted by $\underline{d}(\mathcal{R}_{\mathcal{B}_i}, T)$, is $\min\{d(r^s \in \mathcal{R}_{\mathcal{B}_i}^s, T)\}$. Now, the upper bound of the cohesiveness of $\mathcal{R}_{\mathcal{B}_i}$, denoted by $\bar{\theta}(\mathcal{R}_{\mathcal{B}_i}, T)$ is computed using $\underline{d}(\mathcal{R}_{\mathcal{B}_i}, T)$ in Eq. 8. The upper bound of the σ of a result r_l of query $q_l \in \mathcal{B}_i$ is:

$$\bar{\sigma}(r_l, q_l, T) = \lambda(q_0, q_l) \times \bar{\theta}(\mathcal{R}_{\mathcal{B}_i}, T) \quad (10)$$

Assume $\mathcal{B}'_i = \{q_1, q_2, \dots, q_{l-1}\}$ and \mathcal{R}^* is the k results of $\mathcal{Q}' = \{\mathcal{B}_1, \dots, \mathcal{B}_{i-1}, \mathcal{B}'_i\}$. Also, σ^{min} is the min-score for \mathcal{R}^* . Then, we can stop processing the candidate queries $\{q_l, \dots, q_{|\mathcal{B}_i|}\} \in \mathcal{B}_i$ if $\bar{\sigma}(r_l, q_l, T) < \sigma^{min}$. We call the above pruning technique as *intra-batch* pruning if $l > 1$, otherwise, we call it *inter-batch* pruning (prune the entire batch).

4.2.2 The Framework

Algorithm 4 presents the framework for **batch query processing** technique (BA-QP) of the no-but-semantic-match problem. Similar to first 6 lines of Algorithm 1, it generates the candidate queries, measures their similarity, sorts them and does the initializations in line 1 to apply *inter-query* pruning first. That is, at each step of iteration, we take the topmost similar query $q' \in \mathcal{Q}$ in line 2 and stop processing the candidate queries in \mathcal{Q} as soon as we find a query $q' \in \mathcal{Q}$ such that $|\mathcal{R}^*| = k$ and $q'.sim < \sigma^{min}$ as given in lines 3-4.

Otherwise, we construct the sub-optimal candidate query batch \mathcal{B}^* for q' as given in lines 5-10. In line 12, the shared part \mathcal{K}^s of the candidate query batch \mathcal{B}^* is processed and its results are stored in \mathcal{R}^s . Then, for each query $q'' \in \mathcal{B}^*$, the upper bound of its actual results' score $\bar{r}.score$ is computed in line 14. In lines 15-16, we apply *inter and intra-batch* pruning where we stop processing the candidate query batch \mathcal{B}^* if $|\mathcal{R}^*| = k$ and $\bar{r}.score < \sigma^{min}$. Otherwise, if $|\mathcal{R}^*| = k$, we update σ^{min}

Algorithm 4: BA-QP

Input : User Query q_0 , Tuning Parameter α ,
Keyword Inverted Lists \mathcal{S}

Output: Top-k Semantically Related Result \mathcal{R}^*

- 1 Repeat Lines 1-6 of Algorithm 1.
- 2 **while** $q' \leftarrow getNext(\mathcal{Q}) \neq null$ **do**
- 3 **if** $\mathcal{R}^* \cdot getSize() = k$ **and** $q' \cdot sim < \sigma^{min}$ **then**
- 4 **break** ; // inter-query pruning
- 5 **for** $i = 1 \rightarrow n$ **do**
- 6 $\mathcal{K}^s \leftarrow q' \setminus k_i$
- 7 $\mathcal{B} \leftarrow getBatch(\mathcal{Q}, \mathcal{K}^s)$
- 8 $\mathcal{K}^u \leftarrow getUnique(\mathcal{B}, \mathcal{K}^s)$
- 9 $\mathcal{B}.add(\mathcal{B}_i, c(\mathcal{K}^s, \mathcal{K}^u))$
- 10 $\mathcal{B}^* \leftarrow getMinCost(\mathcal{B})$; $\mathcal{S}' \leftarrow \{\}$;
- 11 retrieve Inverted Lists $\forall k_i \in \mathcal{B}^*.\mathcal{K}^s$ into \mathcal{S}' ;
- 12 $\mathcal{R}^s \leftarrow sharedPartComputation(\sigma^{min}, \alpha, q'', \mathcal{S}')$
- 13 **while** $q'' \leftarrow getNext(\mathcal{B}^*) \neq null$ **do**
- 14 $\bar{r} \cdot score \leftarrow q'' \cdot sim \times coh(\mathcal{R}^s, T)$;
- 15 **if** $\mathcal{R}^* \cdot getSize() = k$ **and** $\bar{r} \cdot score < \sigma^{min}$
 then
- 16 **break**; // inter and intra-batch pruning
- 17 **if** $\mathcal{R}^* \cdot getSize() = k$ **then**
- 18 $root \leftarrow \mathcal{R}^* \cdot root()$; $\sigma^{min} \leftarrow root \cdot score$;
- 19 $k^u \leftarrow q'' \setminus \mathcal{B}^*.\mathcal{K}^s$;
- 20 $\mathcal{S}_1 \leftarrow retrieveKeywordInvertedList(k^u, \mathcal{S})$;
- 21 $\mathcal{R}^* \leftarrow mergeResults(\mathcal{R}^*, \mathcal{R}^s, \sigma^{min}, \alpha, q'', \mathcal{S}_1)$;
- 22 $\mathcal{Q} \leftarrow \mathcal{Q} \setminus \mathcal{B}^*$
- 23 **return** \mathcal{R}^*

by reading the root entry of the heap \mathcal{R}^* in lines 17-18. In lines 19-20, the unshared keyword part k^u of the batch query q'' and the corresponding inverted list \mathcal{S}_1 are retrieved. The procedure *mergeResults* in line 21 generates the final results for q'' by merging shared part results \mathcal{R}^s with the unshared part k^u and inserts the promising results into \mathcal{R}^* . Finally, the queries $q'' \in \mathcal{B}^*$ are excluded from the \mathcal{Q} as pseudocoded in line 22. The above steps continue until \mathcal{Q} becomes empty.

The details of the *sharedPartComputation* method in line 12 of Algorithm 4 is presented in Algorithm 5. In lines 2-5, we compute the result root v_{slca} . Then, we compute all the ancestors of v_{slca} which is implemented in the function *getAncestors* and put them in the potential result roots \mathcal{A} in line 6. For each potential result root $v_{slca}^s \in \mathcal{A}$, we check if it is already processed and is in memory \mathcal{R}^s in line 8. If $v_{slca}^s \notin \mathcal{R}^s$, this result has not been processed and thus, we compute the closest nodes under v_{slca}^s in the inverted lists $\mathcal{S}_i, \forall i \in [1, n-1]$, their distance to v_{slca} , and their score (lines 9-12). In lines 13-14, we check if the score can beat σ^{min} (intra-query pruning). In lines 15-16, r is updated and inserted into \mathcal{R}^s because it is promising. In line 17, we update the anchor match node m_a for the next iteration.

The details of the *mergeResults* method in line 21 of Algorithm 4 is presented in Algorithm 6. In lines 2-5, we compute the potential result root v_{slca}^u which is yet to be confirmed. If the previous result root v_{slca} is

Algorithm 5: BA-QP (sharedPartComputation)

Input : $\sigma^{min}, \alpha, q', \mathcal{S}'$

Output: Shared Result \mathcal{R}^s

- 1 $\mathcal{R}^s \leftarrow null$;
- 2 $m_a \leftarrow getAnchor(\{getNext(\mathcal{S}_i), \forall i \in [1, n-1]\})$;
- 3 **while** $m_a \neq null$ **do**
- 4 $m_i \leftarrow closest(m_a, \mathcal{S}_i), \forall i \in [1, n-1]$ and $i \neq a$;
- 5 $v_{slca} \leftarrow lca(m_1, \dots, m_{n-1})$;
- 6 $\mathcal{A} \leftarrow v_{slca} \cup getAncestors(v_{slca})$;
- 7 **while** $v_{slca}^s \leftarrow getNext(\mathcal{A}) \neq null$ **do**
- 8 **if** $v_{slca}^s \notin \mathcal{R}^s$ **then**
- 9 **for** $i = 1 \rightarrow n-1$ **do**
- 10 $m_i^{li} \leftarrow getTight(\mathcal{S}_i, \mathcal{S}_i.cursor)$;
- 11 $d_i \leftarrow getDist(v_{slca}^s, m_i^{li})$; $r.d \leftarrow r.d + d_i$;
- 12 $r \cdot score \leftarrow q' \cdot sim \times \frac{1}{\log_\alpha(r.d+1)+1}$;
- 13 **if** $r \cdot score < \sigma^{min}$ **then**
- 14 stop reading lists, jump to line 6.
- 15 $r \leftarrow (v_{slca}^s, \{m_i^{li}, \forall i \in [1, n]\})$;
- 16 $\mathcal{R}^s \cdot insert(r)$;
- 17 $m_a \leftarrow getAnchor(\{getNext(\mathcal{S}_i), \forall i \in [1, n-1]\})$;
- 18 **return** \mathcal{R}^s

not an ancestor of v_{slca}^u , then v_{slca} is confirmed as the SLCA result root in line 6. Thus, from \mathcal{R}^s , we retrieve the precomputed closest match nodes $m_i^{li}, \forall i \in [1, n-1]$ of the shared part which is implemented in function *retrieveNode* in line 8 and the distance of the shared part which is implemented in the function *retrieveDist* in line 9. Then, in lines 10-11 we compute the closest match node m_n^{ln} for the unshared part and its distance to v_{slca} and finally, add this distance to the total distance of r . In lines 13-15, we compute $r.score$ and if $r.score > \sigma^{min}$, the result is added to \mathcal{R}^* . We update v_{slca} with the current result root v_{slca}^u if $v_{slca}^u \not\prec_a v_{slca}$ in lines 16-17. Then, the anchor match node m_a is updated in line 18. In lines 19-21, if the last result root is SLCA and is not considered, similar steps are taken and the corresponding result is inserted into \mathcal{R}^* if it is promising.

5 Experiments

This section evaluates the effectiveness and the efficiency of our approach for solving no-but-semantic-match problem in XML keyword search. To compare our results, we adapt and implement the closely related existing XOntoRank method [34] which uses ontology to enhance the XML keyword search on medical datasets. The adaptation is achieved as follows. We create a hash map for the ontologically relevant keywords to the original keywords by using ontological knowledge base. Then for each keyword $k_i \in q_0$ we look up the hash map and find the candidates and put them in the set of candidate keywords \mathcal{K} . Then we create Onto-DIL for each

Algorithm 6: BA-QP (mergeResults)

Input : $\mathcal{R}^*, \mathcal{R}^s, \sigma^{min}, \alpha, q'', S_1$
Output: Top-k Result \mathcal{R}^*

```

1  $r \leftarrow null; v_{slca} \leftarrow null;$ 
2  $m_a \leftarrow getAnchor(\{getNext(S_1), getNext(S_2)\});$ 
3 while  $m_a \neq null$  do
4    $m_i \leftarrow closest(m_a, S_i), \forall i \in [1, 2]$  and  $i \neq a;$ 
5    $v_{slca}^u \leftarrow lca(m_1, m_2);$ 
6   if  $r \neq null$  and  $v_{slca} \not\prec_a v_{slca}^u$  then
7     if  $v_{slca} \in \mathcal{R}^s$  then
8        $m_i^{li} \leftarrow retrieveNode(r, \mathcal{R}^s), \forall i \in [1, n-1];$ 
9        $d \leftarrow retrieveDist(r, \mathcal{R}^s);$ 
10       $m_n^{ln} \leftarrow getTight(S_1, r.cursor_1);$ 
11       $d \leftarrow d + getDist(v_{slca}, m_n^{ln});$ 
12       $r \leftarrow (v_{slca}, \{m_i^{li}, \forall i \in [1, n]\});$ 
13       $r \cdot score \leftarrow q'' \cdot sim \times \frac{1}{\log_\alpha(d+1)+1};$ 
14      if  $r \cdot score > \sigma^{min}$  then
15        | update top-k list  $\mathcal{R}^*$  with  $r;$ 
16      else if  $v_{slca}^u \not\prec_a v_{slca}$  then
17        |  $v_{slca} \leftarrow v_{slca}^u; r.add(S_1.cursor);$ 
18       $m_a \leftarrow getAnchor(\{getNext(S_i), \forall i \in [1, 2]\});$ 
19 if  $v_{slca}^u \not\prec_a v_{slca}$  then
20   | make  $r$  by repeating lines 8-12;
21   | score  $r$  and insert to  $\mathcal{R}^*$  if promising;
22 return  $\mathcal{R}^s$ 

```

Algorithm 7: XO-QP

Input : User Query q_0 , Data T , Ontology O
Output: Top-k Semantically Related Results \mathcal{R}^*

```

1  $H \leftarrow createHashMap(q_0, O);$ 
2  $\sigma^{min} \leftarrow 0;$ 
3 foreach  $k_i \in q_0$  do
4    $\mathcal{K} \leftarrow findRelevant(k_i, H);$ 
5    $S' \leftarrow createDIL(k_i \cup \mathcal{K}, T);$ 
6    $S'.add(NS(k_i, k)), \forall k \in \mathcal{K};$ 
7    $S_{onto} \leftarrow S_{onto} \cup S';$ 
8 while inverted lists in  $S_{onto} \neq null$  do
9    $r \leftarrow computeLCA(S_{onto});$ 
10   $r.score \leftarrow \sum m.NS, \forall m \in r;$ 
11  if  $r.score > \sigma^{min}$  then
12    | update top-k list  $\mathcal{R}^*$  with  $r;$ 
13    |  $update(\sigma^{min});$ 
14 return  $\mathcal{R}^*$ 

```

keyword $k_i \in q_0$ and all of its associated candidates in \mathcal{K} . Afterward, we compute the node score for each entry based on the relevance degree of the original keyword k_i to the candidate keywords in \mathcal{K} . Finally, the inverted list is added to Onto-DIL S_{onto} . After creating S_{onto} , we compute LCAs by using Onto-DIL index and for each result r , the score is computed using the node score of its matched nodes. If the result score $r.score$ is better than the threshold σ^{min} , the result is added to top-k list \mathcal{R}^* . The above is pseudocoded in Algorithm 7. We term this method as XO-QP in this paper.

5.1 Settings

Datasets and Queries: We evaluate our algorithms on two real datasets: (a) IMDB 170MB, that includes around 150,000 recent movies and TV series. (b) DBLP 650MB, which contains publications in major journals and proceedings. We use a wide range of queries to test the efficiency of our proposed methods for each dataset. For each test query, we choose keywords which satisfy the followings: (a) a keyword should be used often by the users and (b) a non-existing keyword should have some semantic counterparts, which have direct mapping in the data source. The test queries have no-but-semantic-match problem. Table 2 presents a part of the test queries called sample queries for detailed analysis of efficiency and effectiveness of our methods.

Environment: All algorithms are implemented in C# and the experiments are conducted on a PC with 3.2 GHz CPU, 8 GB memory running 64-bit windows 7.

5.2 Effectiveness

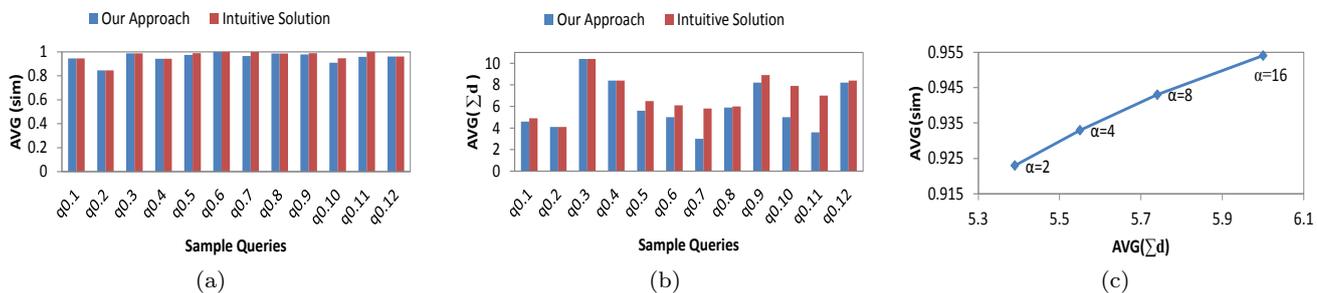
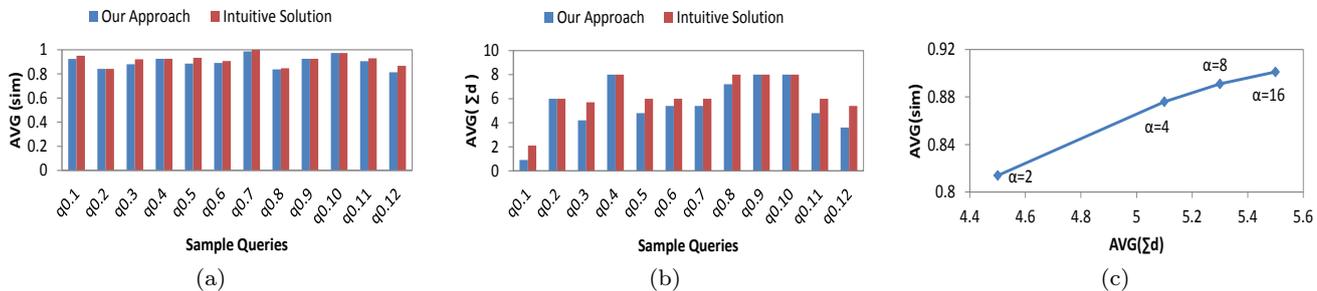
This section evaluates the effectiveness of our approach from different perspectives.

5.2.1 Our Approach versus Intuitive Solution

When a user issues a query and faces an empty result, she might try to change the initial query to obtain some results. However, being unfamiliar with the data source, the user is likely to think of a few synonyms of the keywords of the initial query. Assume the user is a kind of expert and succeeds in constructing the 10 top-most similar queries when changing the initial query. These queries may not produce good results but we consider this *intuitive solution* as a benchmark to gauge the effectiveness of the technique we propose. Here, we compare the top-10 results retrieved by our approach which processes all possible candidate queries with the *intuitive solution* which processes only the 10 topmost similar queries for all test queries given in Table 1. In Fig. 7 (a) and Fig. 8 (a), the average candidate query similarity of the top-10 results for the two approaches are presented for each dataset. Clearly for all the test queries, the average similarity of the results retrieved by our approach is very close to the average candidate query similarity of the results of the *intuitive solution*. Fig. 7 (b) and Fig. 8 (b) show, on average, smaller number of edges (indicate better cohesiveness) for the top-10 results retrieved from our approach compared to the *intuitive solution*. This shows that an ad-hoc approach, even if it is suggested by an expert, is unlikely to find

Table 2: A Sample of test queries for IMDB and DBLP datasets

Test Query Set					
#	IMDB	$ \mathcal{Q} $	#	DBLP	$ \mathcal{Q} $
$q_{0.1}$	ghost, badgering, movie	12	$q_{0.1}$	exigency, analysis, system	252
$q_{0.2}$	battler, spanish, drama	272	$q_{0.2}$	academic, fraudulence, threat	4500
$q_{0.3}$	slump, federal, reserve, harshness, documentary	285	$q_{0.3}$	information, ordination, track	360
$q_{0.4}$	reproach, trespasser, fight, drama	357	$q_{0.4}$	involvement, neuroscience, indicant, information	350
$q_{0.5}$	treasonist, zombie, shiver	1295	$q_{0.5}$	mutter, alarm, analysis	2079
$q_{0.6}$	research, outlander, universe	3528	$q_{0.6}$	deceit, type, analysis	8190
$q_{0.7}$	mass murder, horror, perfidy	11900	$q_{0.7}$	online, trust, selling	276
$q_{0.8}$	victory, exaltation, drama	851	$q_{0.8}$	aftermath, type, analysis, system	225
$q_{0.9}$	partiality, perfidy, fear, english	442	$q_{0.9}$	psychopathy, symptom, visualization, science	3276
$q_{0.10}$	criminal, overcharge, loneliness	12240	$q_{0.10}$	interloper, search, analysis, system	352
$q_{0.11}$	criminal, fear, spanish	1152	$q_{0.11}$	trace, audit, system	238
$q_{0.12}$	victory, exuberance, drama, fight	437	$q_{0.12}$	trace, type, analysis	255

Fig. 7: (a) Average similarity of the candidate queries and (b) average distance of top-10 results for our approach and intuitive solution on IMDB; (c) Average query similarity versus average result distance with varying α .Fig. 8: (a) Average similarity of the candidate queries and (b) average distance of top-10 results for our approach and intuitive solution on DBLP; (c) Average query similarity versus average result distance with varying α .

results that are superior to those provided by the systematic method suggested in this study. In comparison with the *intuitive solution*, our approach also does not retrieve results from the candidate queries that are far from the user given initial query in terms of candidate query similarity. In addition of it, our approach provides the flexibility of trading of the above two aspects, e.g., sometimes users may prioritize cohesive results over similarity to the original query. The tuning parameter α in our approach can be used to balance the weight of these two aspects. Fig. 7(c) and Fig. 8(c), show that as α becomes smaller, the similarity of the results decreases because more priority is given to re-

sult cohesiveness and therefore, the cohesiveness score of the results improves. That is, by setting α to smaller value, our approach effectively retrieves more cohesive results while the similarities of their contributing candidate queries are not very far from the initial query.

5.2.2 Effect of Tuning Parameter

Now, we provide two fine-grained case studies as follows: (1) the influence of the tuning parameter α on the ranking of the retrieved results and (2) trading off between query similarity and result cohesiveness in the final top- k results based on α .

Table 3: Ranking of r_1 is tolerant to α as it is better than r_4 in terms of both query similarity and result cohesiveness

α	$d(r_1, T)$	$\lambda(q_0, q_1) \times \theta(r_1, T)$	$\sigma(r_1, q_1, T)$	$d(r_4, T)$	$\lambda(q_0, q_4) \times \theta(r_4, T)$	$\sigma(r_4, q_4, T)$	$\Delta(r_1, r_4)$
2	7	0.9167×0.25	0.2291	8	0.8462×0.2398	0.2029	0.0262
3	7	0.9167×0.3456	0.3168	8	0.8462×0.3333	0.282	0.0348
4	7	0.9167×0.4	0.3666	8	0.8462×0.3868	0.3273	0.0393
8	7	0.9167×0.5	0.4583	8	0.8462×0.4862	0.4114	0.0469
16	7	0.9167×0.5714	0.5238	8	0.8462×0.5578	0.472	0.0518

Table 4: Trading off candidate query similarity and result cohesiveness in r_2 and r_7

α	$d(r_2, T)$	$\lambda(q_0, q_2) \times \theta(r_2, T)$	$\sigma(r_2, q_2, T)$	$d(r_7, T)$	$\lambda(q_0, q_7) \times \theta(r_7, T)$	$\sigma(r_7, q_7, T)$	$\Delta(r_2, r_7)$
2	11	0.8462×0.2181	0.1845	7	0.7549×0.25	0.1887	-0.0042
3	11	0.8462×0.3065	0.2593	7	0.7549×0.3456	0.2608	-0.0015
4	11	0.8462×0.3581	0.303	7	0.7549×0.4	0.3019	0.0011
8	11	0.8462×0.4555	0.3854	7	0.7549×0.5	0.3774	0.008
16	11	0.8462×0.5273	0.4462	7	0.7549×0.5714	0.4313	0.0149

Case Study-1: This case study demonstrates that our approach is tolerant to the settings of α if one result beats another one in terms of both query similarity and result cohesiveness. This is also expected as the user might explore the top cohesive results with better similarity first. Consider the queries given in Example 7. Here, we extract r_1 from q_1 and r_4 from q_4 as shown in Fig. 4. From Table 3, we see that r_1 will always be ranked better than r_4 as r_1 has the higher overall score σ than r_4 for all settings of α .

Case Study-2: This case study demonstrates how the user can trade off between query similarity and the result cohesiveness based on α . Assume a user would like to explore the results with better cohesiveness first than those with higher similarity. Consider a candidate query $q_7 = \{Jack, academic, position\}$ with $\lambda(q_0, q_7) = 0.7549$ and a result r_7 from q_7 with $d(r_7, T) = 7$. Now, we compare it with r_2 with $d(r_2, T) = 11$ from q_2 with $\lambda(q_0, q_2) = 0.8462$ as given in Example 7. From Table 4, we observe that r_7 outranks r_2 for $\alpha = [2, 3]$. Now, a user needs to set $\alpha > 3$ to explore r_2 before r_7 in the result list by putting more emphasis on query similarity than result cohesiveness.

5.2.3 Our Approach versus XO-QP

In this section, we provide a detailed study of the top candidate query obtained from the top-10 results for different methods. The methods include our proposed method using the tuning parameter $\alpha = 2$ and $\alpha = 16$, intuitive solution (which only executes the top similar queries), XO-QP (which takes the result distance fixed), and XO-QP (which apply the result distance into ranking). In fact, the top-1 query, is the candidate query which has the biggest number of results among top-10 results. The number of results from a candidate query to the total number of top-k results

is defined as $\eta(\mathcal{R}^*, q')$. Therefore, we provide in Table 5 and Table 6 the top-1 candidate queries on IMDB and DBLP datasets respectively. These queries have the maximum $\eta(\mathcal{R}^*, q')$ among the candidate queries which contribute to top-k results. Clearly, our method that set α to 2 has the minimum $d(r, D)$ among other methods in both datasets. That’s because if we set α to a lower number, we put more emphasis on result cohesiveness and therefore, the results with better cohesiveness are ranked higher. However, when we look at our method using $\alpha = 16$, the similarity of the top query in most cases is better comparing to our method using $\alpha = 2$. In the intuitive solution, the top candidate query has the maximum similarity in most cases. Since the intuitive solution only generates top similar queries to retrieve the semantically related results, the similarity of the top candidate query is the maximum in most cases. The XO-QP using fixed distance for ranking the results have the worst cohesiveness score in most cases. That’s because XO-QP uses XRank[13] for retrieving results and therefore, does not limit the results to SLCAs. Moreover, when we do not reflect the distance of the results into the ranking, more sparse results may be ranked higher and inserted to Top- k list. However, when we reflect the distance into ranking of XRank, then the cohesiveness of the results improves.

5.3 Evaluation of Quality

In this section, we evaluate the quality of our approach and compare it with XO-QP. Here, we select some sample queries with no-but-semantic-match problem for both datasets to conduct a comprehensive user study and thereafter, evaluate the overall quality of our approach. In order to carry out a fair user study, we select the users among both experts who have worked in XML

Table 5: Top candidate queries for different methods in IMDB

Original Query	Method	Top-1 Candidate Query	$\eta(\mathcal{R}^*, q')$	$\lambda(q_0, q')$	$d(r \in q', T)$
$q_{0.4}$	Our Approach($\alpha = 2$)	blame, fight, drama, intruder	0.1	0.94	8
$q_{0.4}$	Our Approach($\alpha = 16$)	blame, fight, drama, intruder	0.1	0.94	8
$q_{0.4}$	Intuitive Solution	blame, fight, drama, intruder	0.1	0.94	8
$q_{0.4}$	XO-QP (No Dist)	blame, fight, drama, squatter	0.5	0.92	12
$q_{0.4}$	XO-QP (Dist)	blame, fight, drama, intruder	0.3	0.94	11
$q_{0.7}$	Our Approach($\alpha = 2$)	massacre, panic, betrayal	0.2	0.87	3
$q_{0.7}$	Our Approach($\alpha = 16$)	slaughter, treachery, revulsion	0.1	1	6
$q_{0.7}$	Intuitive Solution	slaughter, treachery, revulsion	0.1	1	6
$q_{0.7}$	XO-QP (No Dist)	uxoricide, apprehension, betrayal	0.1	0.8	8
$q_{0.7}$	XO-QP (Dist)	hit, panic, betrayal	0.3	0.8	3
$q_{0.9}$	Our Approach($\alpha = 2$)	call, betrayal, fear, english	0.1	0.85	4
$q_{0.9}$	Our Approach($\alpha = 16$)	fancy, treachery, fear, english	0.1	1	8
$q_{0.9}$	Intuitive Solution	fancy, treachery, fear, english	0.1	1	8
$q_{0.9}$	XO-QP (No Dist)	bias, treachery, fear, english	0.5	0.93	12
$q_{0.9}$	XO-QP (Dist)	tilt, betrayal, fear, english	0.4	0.93	11
$q_{0.10}$	Our Approach($\alpha = 2$)	murderer, extortion, blood	0.3	0.76	3
$q_{0.10}$	Our Approach($\alpha = 16$)	crook, extortion, desolation	0.1	0.94	7
$q_{0.10}$	Intuitive Solution	crook, extortion, desolation	0.1	0.94	7
$q_{0.10}$	XO-QP (No Dist)	outlaw, extortion, desolation	0.3	0.94	9
$q_{0.10}$	XO-QP (Dist)	outlaw, extortion, desolation	0.3	0.94	9
$q_{0.12}$	Our Approach($\alpha = 2$)	enthusiasm, drama, fight, triumph	0.1	1	8
$q_{0.12}$	Our Approach($\alpha = 16$)	enthusiasm, drama, fight, triumph	0.1	1	8
$q_{0.12}$	Intuitive Solution	enthusiasm, drama, fight, triumph	0.1	1	8
$q_{0.12}$	XO-QP (No Dist)	madness, drama, fight, triumph	0.6	0.94	12
$q_{0.12}$	XO-QP (Dist)	madness, drama, fight, win	0.4	0.87	10

Table 6: Top candidate queries for different methods in DBLP

Original Query	Method	Top-1 Candidate Query	$\eta(\mathcal{R}^*, q')$	$\lambda(q_0, q')$	$d(r \in q', T)$
$q_{0.1}$	Our Approach($\alpha = 2$)	emergency, research, system	0.4	0.88	3
$q_{0.1}$	Our Approach($\alpha = 16$)	emergency, research, system	0.4	0.88	3
$q_{0.1}$	Intuitive Solution	emergency, research, system	0.4	0.88	3
$q_{0.1}$	XO-QP (No Dist)	pinch, breakdown, system	0.6	0.94	6
$q_{0.1}$	XO-QP (Dist)	pinch, breakdown, system	0.6	0.94	6
$q_{0.7}$	Our Approach($\alpha = 2$)	online, content, exchange	0.1	0.52	0
$q_{0.7}$	Our Approach($\alpha = 16$)	online, faith, merchandising	0.1	1	6
$q_{0.7}$	Intuitive Solution	online, faith, merchandising	0.1	1	6
$q_{0.7}$	XO-QP (No Dist)	online, faith, trading	1	1	6
$q_{0.7}$	XO-QP (Dist)	online, faith, trading	1	1	6
$q_{0.8}$	Our Approach($\alpha = 2$)	impact, order, analysis, system	0.1	0.69	0
$q_{0.8}$	Our Approach($\alpha = 16$)	impact, order, analysis, system	0.1	0.69	0
$q_{0.8}$	Intuitive Solution	wake, variety, analysis, system	0.1	0.94	8
$q_{0.8}$	XO-QP (No Dist)	wake, variety, analysis, system	1	0.94	8
$q_{0.8}$	XO-QP (Dist)	wake, variety, analysis, system	1	0.94	8
$q_{0.10}$	Our Approach($\alpha = 2$)	whole, study, analysis, system	0.1	0.24	0
$q_{0.10}$	Our Approach($\alpha = 16$)	intruder, hunt, analysis, system	0.1	1	8
$q_{0.10}$	Intuitive Solution	intruder, hunt, analysis, system	0.1	1	8
$q_{0.10}$	XO-QP (No Dist)	intruder, lookup, analysis, system	0.3	1	8
$q_{0.10}$	XO-QP (Dist)	intruder, lookup, analysis, system	0.3	1	8
$q_{0.11}$	Our Approach($\alpha = 2$)	gesture, testing, system	0.1	0.72	0
$q_{0.11}$	Our Approach($\alpha = 16$)	gesture, testing, system	0.1	0.72	0
$q_{0.11}$	Intuitive Solution	hint, testing, system	0.1	0.9	6
$q_{0.11}$	XO-QP (No Dist)	hint, search, system	0.7	0.9	6
$q_{0.11}$	XO-QP (Dist)	hint, search, system	0.7	0.9	6

keyword search areas and naive users who are graduate computer science students. To do the study, we present to the users with the original queries and their top-10 candidate queries/results retrieved from the top result list \mathcal{R}^* . After that, we ask the users to assess the quality of each candidate query with regards to their semantic similarity to the original query by scoring the candidate queries/results using Cumulated Gain metric [39]. They score each candidate query/result from 0 to 5 points (5 means the best and 0 means the worst).

5.3.1 Ranking Scheme Comparison

The average quality scores of the top-5 and top-10 queries/results for our approach and the existing counterpart XO-QP are presented in Fig. 9 and Fig.10 respectively. From Fig. 9, we observe that our proposed method suggests reasonable results for the no-match query for both top-5 and top-10 results. Also, we see that the average quality of top-5 results are always better than the average quality of top-10 results which indicates that our ranking function successfully ranks more similar and meaningful results higher than the rest of the results. However, we observe that in many cases for the existing method XO-QP, the average quality of results for top-10 is higher than top-5 results (for instance for $q_{0.2}, q_{0.4}, q_{0.11}$ in IMDB and for $q_{0.3}, q_{0.7}, q_{0.10}$ in DBLP) as shown in Fig. 10. This indicates that XO-QP sometimes ranks some less similar and meaningful results higher. This is probably because of adapting XRank scheme into XO-QP. We conclude that our approach addresses the no-but-semantic-match problem better than the existing XO-QP method.

5.3.2 Precision

In this section, we compare the precision of our proposed method with the XO-QP method. In order to make a comparison, we count the number of meaningful results in top-10 results. A result is regarded as meaningful if the average quality score of the assessors is not less than 3. We see that the precision of our approach and XO-QP is presented on both IMDB and DBLP as shown in Fig. 11. In IMDB, we see that the precision of our proposed method is better than XO-QP. This is because we use both semantic similarity and cohesiveness scores to rank the results and retrieve the tightest SLCA results with the maximum similarity to the original query keywords. Also in DBLP, we observe that the precision of our method is better than XO-QP in many cases specially in $q_{0.1}, q_{0.3}$ and $q_{0.8}$. That's because our method can effectively retrieve the most similar results that are cohesive and make a meaningful

combination in terms of data cohesiveness. However, in some cases like $q_{0.11}$ and $q_{0.12}$, we see that the precision of our method is smaller than XO-QP because it uses less similar keywords in the candidate query to retrieve more cohesive results. But the precision is not largely deteriorated as we see in XO-QP in many cases.

5.4 Efficiency

To demonstrate the efficiency of our approach, we use a baseline method which applies inter query pruning only, and uses scan eager for query processing. We call this baseline query processing as BL-QP. The purpose of this baseline is to demonstrate the efficiency of the intra-query pruning scheme in our proposed methods. Overall, we compare the performance of the following methods: (a)BL-QP, (b) SE-QP, (c) AN-QP, and (d) BA-QP, and (e) XO-QP.

5.4.1 Processing Time

Fig. 12 shows the response time of computing the top-10 semantically related results for the sample queries (presented in Table 2) on the IMDB and DBLP datasets. According to the results, BA-QP achieves the best performance. The reason behind this is that we apply inter and intra batch prunings in BA-QP in addition of inter and intra-query prunings. Also, we share the partial results of the shared keywords among the candidate queries in a batch. On average, BA-QP consumed 30 percent time of the time needed by the SE-QP and AN-QP methods. The difference, however, depends on the number of candidate queries. In Fig. 12 (a), the response time for processing the test query $q_{0.1}$ in BA-QP is very close to those of SE-QP and AN-QP due to the small number of candidate queries for $q_{0.1}$ which is 12 and insufficient number of batches (in this case we have only one batch) to apply inter-batch pruning. Moreover, if the cost of merging the shared part \mathcal{K}^s with the unshared part \mathcal{K}^u is relatively high in most of the batches, BA-QP may not outperform the other methods as we see in $q_{0.5}$ on DBLP. In IMDB $q_{0.5}$, however, AN-QP outperforms SE-QP by a large margin because the data distributions in most of the inverted lists are skewed. In such case, many nodes are skipped in AN-QP, which improves the performance. Clearly, BL-QP has the worst performance on most cases because it does not apply the intra query pruning, therefore, incurs unnecessary computations. The XO-QP, however, is comparable to BA-QP in many cases because we set the node score threshold to 0.9 which means that only the relevant nodes to the keywords with the similarity degree bigger than 0.9 are selected. If we set this value

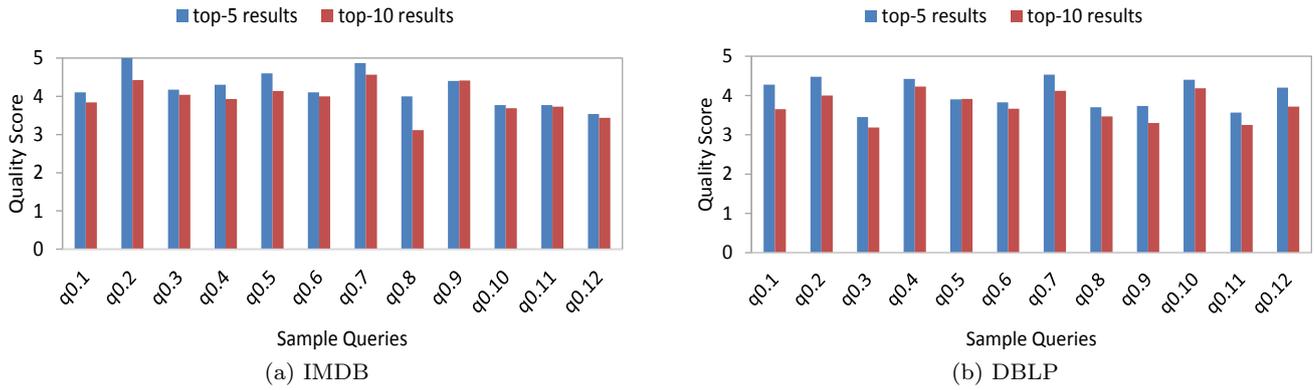


Fig. 9: Average quality of results in our approach: top-5 vs. top-10.

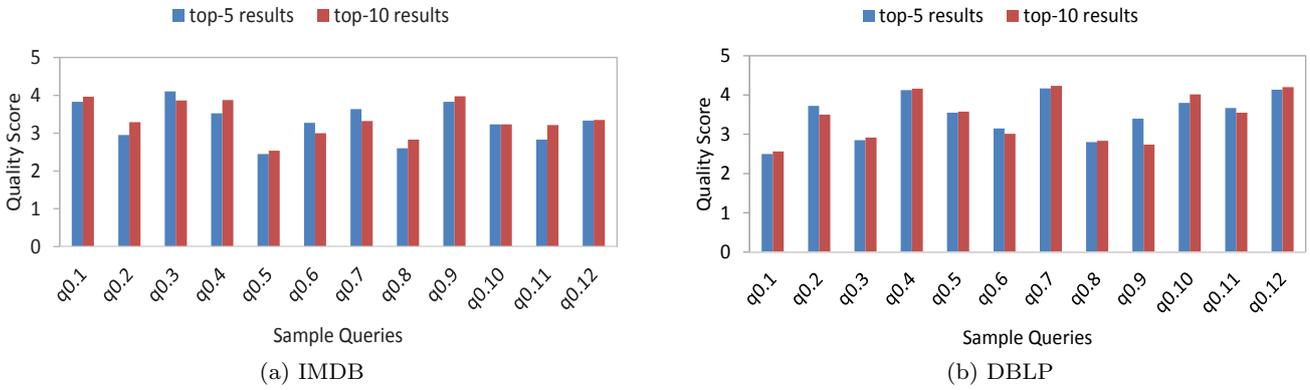


Fig. 10: Average quality of results in XO-QP: top-5 vs. top-10.

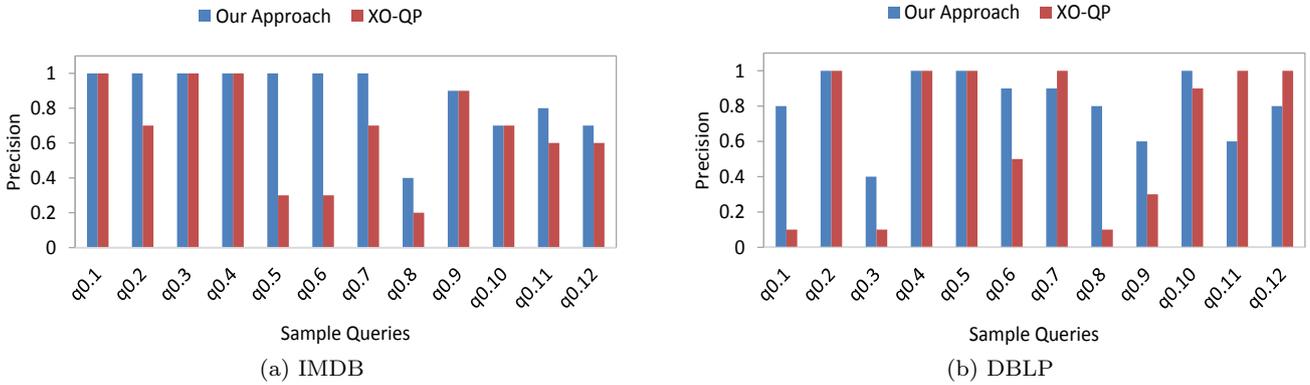
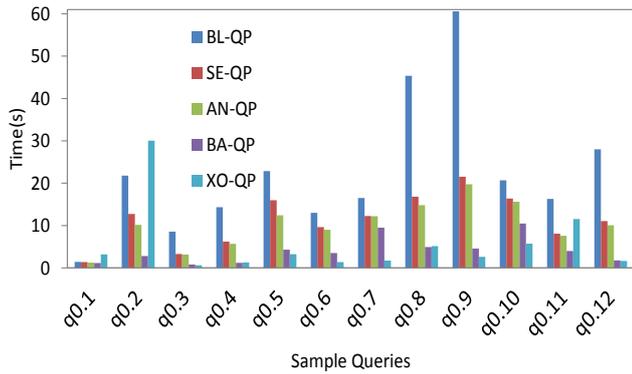


Fig. 11: Precision comparison for computing top-10 results.

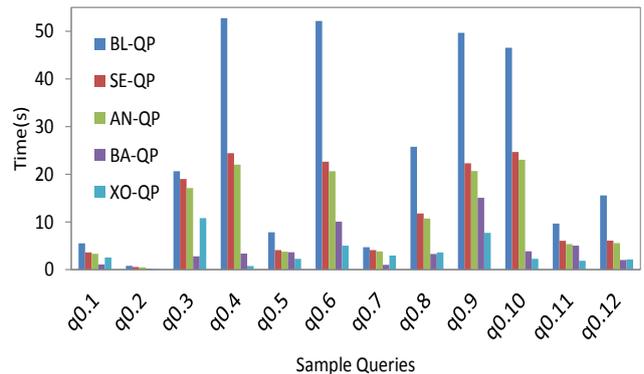
to a smaller number, XO-QP performance deteriorates and gets closer to AN-QP and SE-QP. Moreover, XO-QP builds a special inverted list called XOnto-DIL for the keywords which takes additional time and space for its creation while our proposed methods do not use such indexing and use the normal DIL for query processing.

In Fig.13, the average processing time of different methods for all test queries are presented. Clearly the

baseline method which does not use intra query pruning spends the maximum time for processing on both datasets. The SE-QP and AN-QP spend less time comparing to BL-QP because we apply inter and intra query pruning and therefore, the processing terminates early when we reach to the global σ^{min} . Moreover, there is a narrow improvement in the AN-QP over SE-QP due to using the anchor node processing which skips



(a) IMDB



(b) DBLP

Fig. 12: Processing time of the sample queries for computing top-10 results.

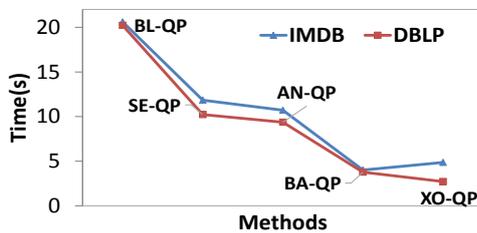
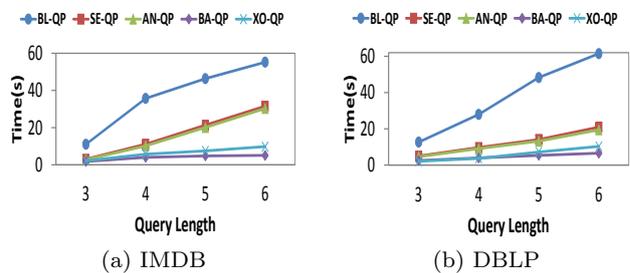


Fig. 13: Average processing time on all test queries

many redundant computations and expedites the efficiency. The BA-QP processing improves sharply on both datasets because in BA-QP, we execute queries in batch and share the computations among the queries in the batch and thereafter, we reach to the global σ^{min} before SE-QP and AN-QP can do. Also, we can apply inter and intra-batch pruning which expedite its efficiency. In XO-QP, however, the performance is close to BA-QP on both datasets because we build the XOntoDIL on the nodes with the relevance degree no less than 0.9. In such condition, only the highly relevant nodes are selected to replace the keywords, therefore, the processing time does not grow considerably due to small number of candidate keywords. In contrast, by setting the threshold to lower numbers, the processing time will grow exponentially and gets closer to SE-QP. Furthermore, our methods do not need to build special indexing, therefore, avoid using additional space and offline processing time for building such indexing.

5.4.2 Effect of Query Length

In this experiment, we choose test queries that have at least 100 candidate queries. At each step, we set their length to a number of settings ($\{3, 4, 5, 6\}$) and compute the average response time of all queries when processing 100 candidate queries while setting k to 10. In each step, for the test queries that have smaller number of



(a) IMDB

(b) DBLP

Fig. 14: Average processing time of the test queries with varying query length (i.e., number of keywords).

keywords, we add additional keywords to increase their length. We also use the same queries to conduct experiments and compare the results with XO-QP method. In Fig. 14 the effect of growing the length of queries on the response time is analyzed. The response time of BA-QP is almost fixed compared to XO-QP, BL-QP, SE-QP, and AN-QP on IMDB. Similarly in DBLP, BA-QP has the slowest growth in response time when the number of keywords increases while BL-QP, SE-QP and AN-QP show a big jump in processing times. Also, XO-QP shows a jump in the processing time when the query length increases. Thus, the performance of XO-QP is sensitive to the query length parameter. The sharpest increase in the processing time is related to BL-QP because it does not apply intra-query pruning, therefore, all the inverted lists are accessed during the processing and this incurs many useless computations specifically when the query length increases.

In summary, we can conclude that BA-QP method is not sensitive to the number of keywords due to sharing the computations while the performance of SE-QP and AN-QP is highly sensitive to the query length.

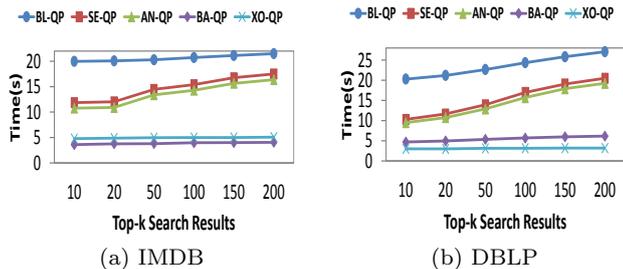


Fig. 15: Average processing time of the test queries with varying k in top- k list.

5.4.3 Effect of Top- k List Size k

In this experiment, we vary the top- k size k and compute the average processing time. In Fig. 15, we observe the effect of top- k size k on the average processing time for different methods. Clearly, the BA-QP and XO-QP methods are not that sensitive to the value of k . By increasing the value of k , processing times of BA-QP and XO-QP show only a small increase or almost fixed. On the other hand, for SE-QP and AN-QP methods, any increase on k , leads to a considerable jump on the average processing time. However, this increase is not big when we change k from 10 to 20. When k is selected as a bigger number, σ^{min} will be smaller and this causes the inter query pruning in SE-QP and AN-QP to be less effective. That is, the application of inter-query pruning is delayed in SE-QP and AN-QP for big k . However in BA-QP, we execute a candidate query batch by sharing the computations among the queries in it and thereafter, we reach to the global σ^{min} before SE-QP and AN-QP can do. Also, we can apply inter and intra batch pruning with this early found global σ^{min} which helps BA-QP to expedite its efficiency.

5.4.4 Effect of Tuning Parameter α

In this experiment, we vary α between a number of settings ($\{2, 4, 8, 16\}$) and compute the average processing time of the test queries when $k = 10$. Fig. 16 presents the effect of choosing different values for α . Larger values of α reduce the sensitivity to data cohesiveness. This usually leads to a more effective intra-query pruning and decreases the processing time. On the contrary, as α decreases, the possibility for the partial results scores to be smaller than the σ^{min} also decreases. In this case, the intra-query pruning becomes less effective and therefore, the processing time increases, specifically, in SE-QP and AN-QP. In BL-QP, however, there is no significant difference in the processing time when α changes to smaller number. This is because BL-QP

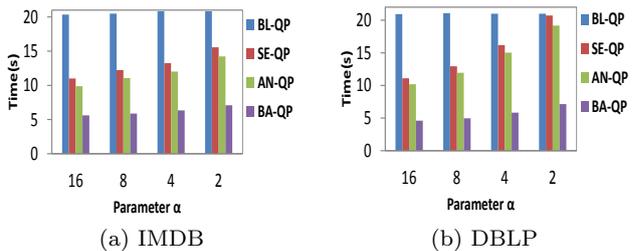


Fig. 16: Average processing time of the test queries with varying tuning parameter α .

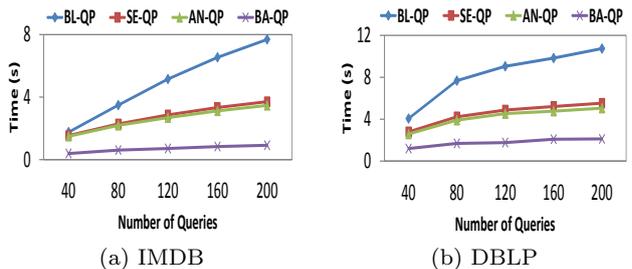


Fig. 17: Processing time with varying number of candidate queries $|\mathcal{Q}|$.

does not apply intra query pruning and therefore, the processing time is not affected that much by α .

5.4.5 Effect of Candidate Queries Number $|\mathcal{Q}|$

For this experiment, we choose test queries which have at least 200 candidate queries. At each step we process a certain number ($\{40, 80, 120, 160, 200\}$) of their candidate queries and compute the average processing time when $k = 10$. In Fig. 17, the effect of growing the number of candidate queries $|\mathcal{Q}|$ on the response time is analyzed. The response time of BA-QP method grows slowly compared to BL-QP, SE-QP, and AN-QP which show a sharp rise in each step. In BL-QP, the growth in the query processing time is the maximum among the methods because it does not apply intra query pruning, therefore, it incurs many unnecessary computations during execution of the candidate queries. We can conclude that BA-QP is not that sensitive to $|\mathcal{Q}|$. This is because BA-QP not only shares the computations among the candidate query batch but also applies inter and intra batch pruning with the early-found global σ^{min} to expedite its performance.

5.4.6 Pruning Improvement

In this section, we show the pruning effect on the processing time. From Fig. 18, the sample queries processing time are shown for 3 scenarios: (a) when there is no pruning for processing the queries, (b) when only the

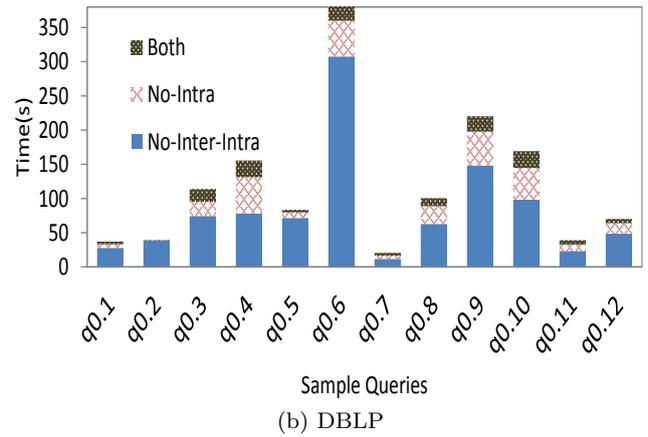
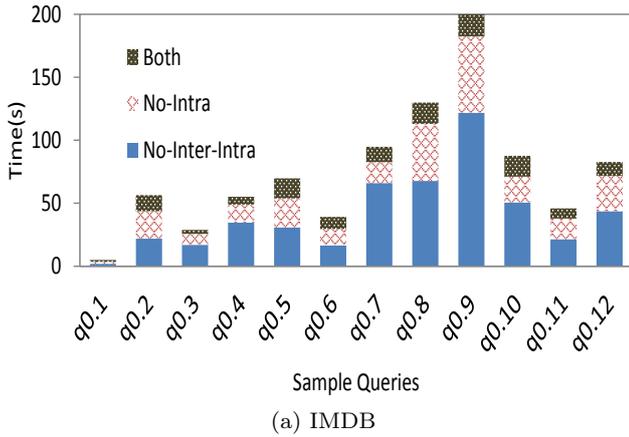


Fig. 18: Inter and Intra Query Pruning Improvement on Sample Queries.

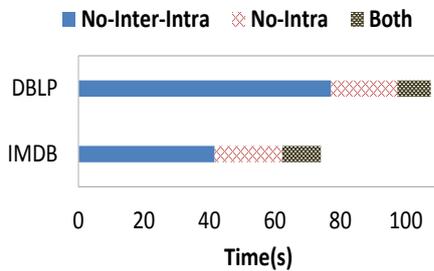


Fig. 19: Average Pruning Improvement.

inter query pruning is implemented, and (c) when both inter and intra pruning methods are implemented. The processing time are measured using SE-QP method. Clearly, the inter query pruning shows the most effective method to cut the processing time on most of the sample queries. If the number of candidate queries is large and the breaking point occurs when most of the queries are not executed, then inter query has the best performance. e.g., in the sample queries $q_{0.9}$ and $q_{0.10}$ on IMDB, the number of executed queries are $\frac{225}{442}$ and $\frac{4918}{12240}$ respectively and in the sample queries $q_{0.6}$ and $q_{0.9}$ on DBLP, the number of executed queries are $\frac{912}{8190}$ and $\frac{1287}{3276}$ respectively. Therefore in these cases, most of the queries are not executed by using inter query pruning and the processing time reduced sharply. The intra query pruning is more effective when the number of query keywords is bigger or the inverted list that is not accessed due to pruning is big sized. In such condition, some inverted lists are not accessed when the result is not able to beat the σ^{min} and this expedites the processing time. For example, in the sample queries $q_{0.8}$ and $q_{0.9}$ on IMDB and $q_{0.4}$ and $q_{0.9}$ on DBLP, the query keywords are from 4 to 5 keywords and include some big sized inverted lists that are not accessed, therefore the processing time reduced considerably.

Fig. 19 presents the average processing time for the set of test queries for 3 scenarios: (a) no pruning is implemented, (b) only inter query pruning is implemented, and (c) both pruning techniques are implemented. Clearly, the processing time for the case with no pruning is maximum on both datasets. This shows that inter query pruning has the most tangible effect on the processing time by avoiding to execute the queries that cannot contribute to the \mathcal{R}^* . The efficiency improvement on IMDB and DBLP is 2 and 3 times respectively. After that, intra query pruning expedites the efficiency by avoiding to access all inverted lists when the result cannot beat σ^{min} .

Fig. 20 presents the sample queries processing time for BA-QP in 2 scenarios: (a) when no batch pruning is implemented, (b) when batch pruning is implemented. In most cases, the processing time for the method which uses pruning has decreased. The batch pruning becomes more effective when the query length increases as shown in $q_{0.3}$ and $q_{0.12}$ on IMDB or in $q_{0.4}$ and $q_{0.8}$ on DBLP. Moreover, when the candidate queries contain some big inverted lists and we reach to the global σ^{min} early, the improvement is more considerable as in $q_{0.1}$ and $q_{0.2}$ on IMDB and in $q_{0.1}$ and $q_{0.7}$ on DBLP.

Fig. 21 shows the average processing time of the test queries for BA-QP on 2 cases: (a) when no batch pruning is implemented, (b) when batch pruning is implemented. We observe from the picture that the batch pruning improved the processing time on both datasets, however, the improvement on DBLP is more considerable. The improvement is achieved because we reach to the global σ^{min} earlier by applying batch pruning, therefore it expedites the performance.

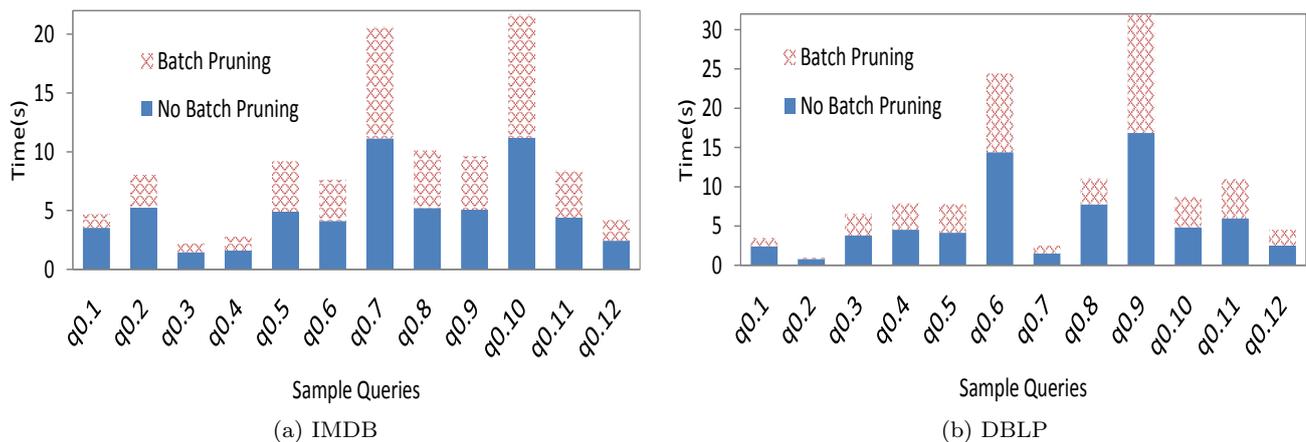


Fig. 20: Batch Pruning Improvement on Sample Queries.

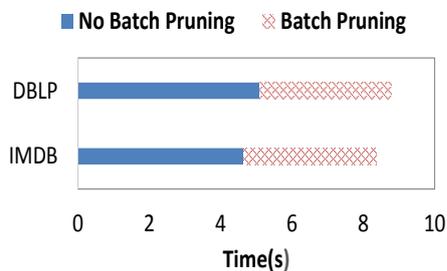


Fig. 21: Average Batch Pruning Improvement.

6 Related Work

Failed Queries. When a user queries a data source, the result may be empty or otherwise below expectation. This problem known as failed queries has inspired a broad range of research in the database community (e.g. [30],[31],[32]). In the context of relational databases, the problem has been studied by Nambiar and Kambhampati [21], Muslea [19] as well as Muslea and Lee [20]. Nambiar and Kambhampati [21] presented approximate functional dependencies to relax the user original query and find tuples similar to the user query. Muslea [19], as well as Muslea and Lee [20] used machine learning techniques to infer rules for generating replacement queries. Amer-Yahia, Cho and Srivastava [3], Brodianskiy and Cohen [6] as well as Cohen and Brodianskiy [9] studied query relaxation for XML data. The studies proposed to discover the constraints in the queries that prevent results from being generated and remove them so that a result can be produced. All these investigations focus on the modification of the user original query constraints on the content level rather than the semantic analysis of the original query constraints. Hill et al. [29], used the ontology information to relax structured XML queries. Farfan et al. [34], proposed XOntoRank system to address the ontology-aware XML key-

word search of electronic medical records. Unlike their work which uses SNOMED ontology for enhancing the search on medical records, we address the general no-match problem on XML data and use a general ontological knowledge base like a thesaurus or dictionary to solve the problem for general documents.

Query Expansion. Query expansion has widely studied in many works such as [35],[36],[37],[38]. Schenkel, Theobald, and Weikum [36] proposed XXL which combines the keyword search with structural conditions and semantic similarity to increase the quality of results. Kim and Kong [37] suggested a query expansion technique that uses an ontology algorithm to map a target DTD to ontology. This scheme is successful for expanding the queries minimally. Kim, Kong, and Jeon [38] developed a web XML document search engine that applies ontology-DTD match algorithm for remote documents. However, in all of the above works, the focus is on structured queries. In our work, we find some semantic counterparts for specific non-mapped keywords for replacement, therefore, query expansion is not useful in our case.

Recommendation Systems. Users are often interested in items similar to those they have visited before or to content that has been looked up by similar users. These items are presented by the recommendation systems. Akbarnejad et al. [1] and Chatzopoulou, Eirinaki and [8] proposed query recommendation based on a prediction of the items that user is interested in. Yao et al. [27] proposed to exploit structural semantics for query reformulation. Meng, Cao and Shao [17] used the semantic relationships between keywords and keyword queries to suggest a set of keyword queries from the query log. However, the semantic relationship is interpreted as the co-occurrence of the keywords and no ontological analysis is carried out. Moreover, the work focuses on extracting similar queries from the

query log using data mining techniques without processing the results. Drosou and Pitoura [11] presented a database exploration framework which recommends additional items called “You May Also Like” results. However, the recommended results are compiled based on the results of the original query and there is no focus on semantic connection between the original query and recommended results.

Mismatch Problem. Sometimes the system shows erroneous mismatch results for a user query which is called mismatch problem. Bao et al. [4] proposed a framework to detect the keyword queries that lead to a list of irrelevant results on XML data. They detect a mismatch problem by analyzing the results of a user query and inferring the user’s intended node type result based on data structure. Based on this, they are able to suggest queries with relevant results to the user. Unlike the current study, Bao et al. investigate ways of producing relevant results instead of finding results for no-match queries.

Query Cleaning. Sometimes the empty result is caused by typographical errors. Pu and Yu [22] and Lu et al. [16] investigated a way of suggesting queries that have been cleaned of typing errors. Unlike our study, these authors do not tackle the problem of non-mapped keywords.

Ontology-based Querying. Many studies have used ontology information for searching the semantic web [10], [15], [33]. Studies by Aleman-Meza [2], Cakmak and Özsoyoglu [7] as well as Wu, Yang and Yan [24] used ontology information to find frequent patterns in graphs. Wu, Yang and Yan [24] proposed an improved subgraph querying technique by ontology information. They revised subgraph isomorphism by mapping a query to semantically related subgraphs in terms of a given ontology graph. Our work generates substitute queries for the user given keyword query by extracting the semantically related keywords from the ontological knowledge base and thereafter, produce semantically related results to the user query instead of returning an empty result set to the user.

7 Conclusion

This paper investigates ways of efficiently building substitute queries against XML data sources when the user given keyword query fails to produce any result as one or more of its keywords do not exist in the data source. Our approach depends on an ontological knowledge base for a discovery of semantically related keywords to generate the substitute queries, which can be executed against the data source to produce the semantically related results for the user’s original query. As the number of

substitute queries can be potentially large and also, not all semantically related results are meaningful to the same degree, we propose efficient pruning techniques to reduce the number of substitute queries and return only the top- k semantically related results. We develop two query processing algorithms to evaluate the substitute queries against the data source based on our pruning techniques. We also develop a batch processing technique that exploits the shared keywords among the substitute queries to expedite the performance further. The extensive experiments with two real datasets validate the effectiveness and efficiency of our approach.

8 Acknowledgment

This work is supported by the Australian Research Council discovery grants DP140103499 and DP160102412.

References

1. J. Akbarnejad, G. Chatzopoulou, M. Eirinaki, S. Koshy, S. Mittal, D. On, N. Polyzotis, and J. S. V. Varman. SQL query recommendations. *PVLDB*, 3(2):1597–1600, 2010.
2. B. Aleman-Meza, C. Halaschek-Wiener, S. S. Sahoo, A. P. Sheth, and I. B. Arpinar. Template based semantic similarity for security applications. In *ISI*, pages 621–622, 2005.
3. S. Amer-Yahia, S. Cho, and D. Srivastava. Tree pattern relaxation. In *EDBT*, pages 496–513, London, UK, UK, 2002. Springer-Verlag.
4. Z. Bao, Y. Zeng, T. W. Ling, D. Zhang, G. Li, and H. V. Jagadish. A general framework to resolve the mismatch problem in XML keyword search. *VLDB J.*, 24(4):493–518, 2015.
5. P. Bouquet, G. M. Kuper, and S. Zanobini. Asking and answering queries semantically. In *WOA*, pages 22–27, 2005.
6. T. Brodianskiy and S. Cohen. Self-correcting queries for xml. In *CIKM*, pages 11–20, New York, NY, USA, 2007. ACM.
7. A. Cakmak and G. Özsoyoglu. Taxonomy-superimposed graph mining. In *EDBT*, pages 217–228, 2008.
8. G. Chatzopoulou, M. Eirinaki, and N. Polyzotis. Query recommendations for interactive database exploration. In *SSDBM*, pages 3–18, 2009.
9. S. Cohen and T. Brodianskiy. Correcting queries for xml. *Inf. Syst.*, 34(8):757–777, Dec. 2009.
10. O. Corby, R. Dieng-Kuntz, C. Faron-Zucker, and F. L. Gandon. Searching the semantic web: Approximate query processing based on ontologies. *IEEE Intelligent Systems*, 21(1):20–27, 2006.
11. M. Drosou and E. Pitoura. Ymaldb: exploring relational databases via result-driven recommendations. *VLDB J.*, 22(6):849–874, 2013.
12. J. Feng, G. Li, J. Wang, and L. Zhou. Finding and ranking compact connected trees for effective keyword proximity search in XML documents. *Inf. Syst.*, 35(2):186–203, 2010.
13. L. Guo, F. Shao, C. Botev, and J. Shanmugasundaram. XRank: ranked keyword search over XML documents. In *SIGMOD*, pages 16–27, 2003.

14. V. Hristidis, N. Koudas, Y. Papakonstantinou, and D. Srivastava. Keyword proximity search in XML trees. *IEEE Trans. Knowl. Data Eng.*, 18(4):525–539, 2006.
15. E. Little, K. Sambhoos, and J. Llinas. Enhancing graph matching techniques with ontologies. In *FUSION*, pages 1–8, 2008.
16. Y. Lu, W. Wang, J. Li, and C. Liu. Xclean: Providing valid spelling suggestions for XML keyword queries. In *ICDE*, pages 661–672, 2011.
17. X. Meng, L. Cao, and J. Shao. Semantic approximate keyword query based on keyword and query coupling relationship analysis. In *CIKM*, pages 529–538, 2014.
18. G. A. Miller, R. Beckwith, C. Fellbaum, D. Gross, and K. Miller. Wordnet: An on-line lexical database. *International Journal of Lexicography*, 3:235–244, 1990.
19. I. Muslea. Machine learning for online query relaxation. In *KDD*, pages 246–255, New York, NY, USA, 2004. ACM.
20. I. Muslea and T. J. Lee. Online query relaxation via bayesian causal structures discovery. In *AAAI*, pages 831–836. AAAI Press, 2005.
21. U. Nambiar and S. Kambhampati. Answering imprecise queries over autonomous web databases. In *ICDE*, page 45, 2006.
22. K. Q. Pu and X. Yu. Keyword query cleaning. *PVLDB*, 1(1):909–920, 2008.
23. C. Sun, C. Y. Chan, and A. K. Goenka. Multiway slca-based keyword search in XML data. In *WWW*, pages 1043–1052, 2007.
24. Y. Wu, S. Yang, and X. Yan. Ontology-based subgraph querying. In *ICDE*, pages 697–708, 2013.
25. Z. Wu and M. Palmer. Verbs semantics and lexical selection. In *ACL*, pages 133–138, Stroudsburg, PA, USA, 1994.
26. Y. Xu and Y. Papakonstantinou. Efficient keyword search for smallest leas in XML databases. In *SIGMOD*, pages 537–538, 2005.
27. J. Yao, B. Cui, L. Hua, and Y. Huang. Keyword query reformulation on structured data. In *ICDE*, pages 953–964, 2012.
28. L. Yao, C. Liu, J. Li, and R. Zhou. Efficient computation of multiple XML keyword queries. In *WISE*, pages 368–381, 2013.
29. J. Hill, J. Thorson, B. Guo, and Z. Chen. Toward ontology-guided knowledge-driven XML query relaxation. In *Second International Conference on Computational Intelligence, Modeling and Simulation*, pages 448–453, 2010.
30. W.W. Chu, H. Yang, K. Chiang, M. Minock, G. Chow, and C. Larson CoBase: A scalable and extensible cooperative information system. In *Intelligent Inf. Syst.*, pages 6(2):223–259, 1996.
31. S.Y Huh, K.H. Moon, and J.K. Ahn Cooperative query processing via knowledge abstraction and query relaxation. In *Advanced Topics in Database Research*, pages 1:211–228, 2002.
32. X. Zhou, J. Guagaz, W.T. Balke, and W. Nejdl Query relaxation using malleable schemas. In *SIGMOD*, pages 1:545–556, 2007.
33. J. Mei, L. Ma, and Y. Pan Ontology query answering on databases. In *International Semantic Web Conference*, pages 445–458, 2006.
34. F. Farfan, V. Hristidis, A. Ranganathan, and M. Weiner XOntoRank: Ontology-aware search of electronic medical records. In *ICDE*, pages 820–831, 2009.
35. J. Xu, and W.B. Croft Query expansion using local and global document analysis. In *SIGIR*, pages 4–11, 1996.
36. R. Schenkel, A. Theobald, and G. Weikum Semantic similarity search on semistructured data with the XXL search engine. In *Information Retrieval*, pages 8(4):521–545, 2005.
37. M.S. Kim, and Y.H. Kong Ontology-DTD matching algorithm for efficient XML query. In *Fuzzy Systems and Knowledge Discovery*, pages 3614:1093–1102, 2005.
38. M.S. Kim, Y.H. Kong, and C.W. Jeon Remote-specific XML query mobile agents. In *Data Engineering Issues in E-Commerce and Services*, pages 4055:143–151, 2006.
39. K. Järvelin, J. Kekäläinen Cumulated Gain-based Evaluation of IR Techniques. In *ACM Trans. Inf. Syst.*, pages 20(4):422–446, 2002.