

EASE: An Effective 3-in-1 Keyword Search Method for Unstructured, Semi-structured and Structured Data

Guoliang Li¹ Beng Chin Ooi² Jianhua Feng¹ Jianyong Wang¹ Lizhu Zhou¹

¹Department of Computer Science and Technology, Tsinghua University, Beijing 100084, China

²School of Computing, National University of Singapore, 117543, Singapore

{liguoliang,fengjh,jianyong,dcszlj}@tsinghua.edu.cn; ooibc@comp.nus.edu.sg

ABSTRACT

Conventional keyword search engines are restricted to a given data model and cannot easily adapt to unstructured, semi-structured or structured data. In this paper, we propose an efficient and adaptive keyword search method, called EASE, for indexing and querying large collections of heterogeneous data. To achieve high efficiency in processing keyword queries, we first model unstructured, semi-structured and structured data as graphs, and then summarize the graphs and construct graph indices instead of using traditional inverted indices. We propose an extended inverted index to facilitate keyword-based search, and present a novel ranking mechanism for enhancing search effectiveness. We have conducted an extensive experimental study using real datasets, and the results show that EASE achieves both high search efficiency and high accuracy, and outperforms the existing approaches significantly.

Categories and Subject Descriptors

H.2.8 [Database Applications]: Miscellaneous

General Terms

Algorithms, Performance, Languages

Keywords

Keyword Search, Indexing, Ranking, Graph Index

1. INTRODUCTION

Keyword search is a proven and widely popular mechanism for querying document systems and the World Wide Web. Recently, it has even been extensively applied to extract useful and relevant information from the Internet. Furthermore, the database research community has also recognized the benefits of keyword search and has been introducing keyword search capability into relational databases [1, 2, 6, 13, 15, 20, 22, 23], XML databases [4, 11, 14, 16, 19, 21, 25, 28] and graph databases [10, 12, 17].

However, the existing web search engines cannot integrate information from multiple interrelated pages to answer key-

word queries meaningfully. Next-generation web search engines require *link-awareness*, or more generally, the capability of integrating correlative information items that are linked through hyperlinks. Meanwhile, the efficiency of keyword search on structured and semi-structured data remains a challenging problem. This is so because the traditional approaches have always employed the inverted index to process keyword queries, which is effective for unstructured data but inefficient for semi-structured and structured data. This is because the inverted index is inadequate for identifying the “best” answers with complex structural information, which is rather rich in XML documents or relational databases.

To the best of our knowledge, very few existing studies could be universally applied to unstructured data (e.g., text documents), semi-structured data (e.g., XML documents), structured data (e.g., relational databases) and graph data. Therefore, providing both effective and efficient search ability over such heterogeneous collections within a single search engine remains a big challenge. As it is, the structure of the data, such as the potentially hierarchical embedding in XML documents, is not fully exploited for answering keyword queries. It is also not taken into account for result ranking in most search engines. Consequently, current implementations focus on either IR-style search to meaningfully rank the results but ignore the rich structural information, or DB-style search to discover answers by identifying structural relationships but employ a very straightforward ranking mechanism.

This less-than-ideal situation calls for a framework for indexing and querying over large collections of unstructured, semi-structured or structured data, and adaptive ranking of the results retrieved over those heterogeneous data. In this paper, we propose EASE, an Efficient and Adaptive keyword SEarch method, as an attempt in that direction. Our work is in line with the current trend of seamlessly integrating databases (DB) and information retrieval (IR) techniques [3, 27]. EASE seamlessly integrates efficient query evaluation and adaptive scoring for ranking results. From the DB point of view, EASE provides an efficient algorithmic basis for scalable *top-k*-style processing of large amounts of heterogeneous data for the discovery of rich structural relationships. It works by employing an adaptive, efficient and novel index beyond the inverted index. From the IR viewpoint, EASE integrates an effective ranking mechanism to improve search effectiveness.

In our approach, we model unstructured, semi-structured and structured data as graphs, with nodes being documents, elements and tuples respectively, and edges being hyper-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGMOD '08, June 9–12, 2008, Vancouver, BC, Canada.

Copyright 2008 ACM 978-1-60558-102-6/08/06 ...\$5.00.

links, parent-child relationships (or IDREFS) and primary-foreign-key relationships respectively. We enable efficient keyword queries on these heterogeneous data by summarizing, clustering the graphs and constructing graph indices. To facilitate efficient keyword-based query processing, we examine the issues of indexing and ranking to improve search quality. To the best of our knowledge, this is the first attempt to efficiently and adaptively process keyword queries on such heterogeneous data, and also the first work to propose the novel graph index, which is efficient in identifying rich structural relationships.

Our contributions in this paper are as follows:

- We model unstructured, semi-structured and structured data as graphs and propose an efficient keyword search method, EASE, to adaptively process keyword queries over the heterogeneous data. We devise an effective graph index as opposed to the inverted index, to improve search efficiency and effectiveness.
- We propose a novel ranking mechanism for effective keyword search by taking into account both the structural compactness of answers from the DB viewpoint and the textual relevancy from the IR point of view.
- We examine the issues of indexing and ranking, and devise a simple and yet efficient indexing mechanism to index the structural relationships between the transformed data. The index is amenable to the deployment of existing *top-k* ranking methods.
- We have conducted an extensive performance study using real datasets and various queries with different characteristics. The results show that EASE achieves both high search efficiency and accuracy, and outperforms existing state-of-the-art methods.

The rest of this paper is organized as follows. We present the *r-radius Steiner graph problem* in Section 2. Section 3 introduces a novel graph index. We present a novel scoring function in Section 4. We examine the issues of indexing and ranking, and propose an indexing mechanism in Section 5. Extensive experimental evaluations are provided in Section 6. We review the related work in Section 7 and conclude the paper with Section 8.

2. ADAPTIVE KEYWORD SEARCH MODEL

2.1 Motivation

2.1.1 Unstructured data

Although many prior studies of keyword search over text documents (e.g., HTML documents) have been proposed, they all produce a list of individual pages as results. In the event that there are no pages that contain all the keywords, they will return pages with some of the input keywords ranked by relevancy. Even if two or more interrelated pages contain all the keywords, the existing methods cannot integrate the pages into one relevant and meaningful answer. For example, to search for conferences covering the topic of “Data Integration” held in “Canada” in 2008, one may issue a keyword query of “Conference 2008 Canada Data Integration” to a search engine such as GOOGLE. As we all know, the venue of SIGMOD 2008 is Canada and “Data Integration” is one of its major research topics. Yet surprisingly, the homepage of SIGMOD 2008 is neither in the top 10 results nor even in the first 100 answers. Why? The reason is that SIGMOD

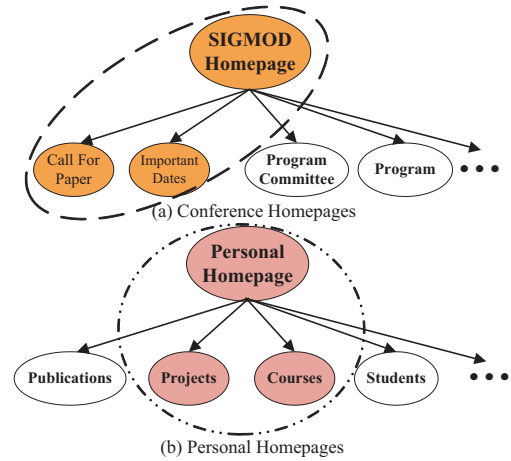


Figure 1: Homepages

2008 splits its information into several pages methodically as shown in Figure 1(a) and IMPORTANT DATE contains the keywords "2008, Conference" while "Data Integration" is contained in the CALL FOR PAPER page. Such data lineage problem also persists in most recently proposed community information management platforms [5].

Consequently, the existing search engines often include a number of false negatives due to the limitation of their models, which take only a list of individual pages as search results but neglect the fact that interrelated pages linked by hyperlinks may be more meaningful. Yet, this is not an *ad hoc* problem but a ubiquitous one over the Internet. As another example, most researchers organize their homepages according to content, as shown in Figure 1(b). Suppose a user searches for professors who teach a specific course and have a specified project, and inputs some keywords. Although there may be no page that contains all input keywords, the page units composed of HOMEPAGE, PROJECT and COURSE of some professors may answer this query meaningfully.

The accuracy and lineage of data have recently received considerable attention, but mainly from a theoretical perspective. While existing studies such as Trio [24] extend conventional data management by incorporating accuracy and lineage as integral components of both data and queries, few works are *link aware* to search text documents over the Internet and take into consideration the fact that multiple interrelated pages linked by hyperlinks may be more meaningful. Indeed, the accuracy problem of traditional search engines is ubiquitous over the Internet. This is so because most web sites organize their data systematically and relevant data may be separated into different pages but linked through hyperlinks. Although Li et al. [18] have proposed a method of retrieving and organizing web pages by “Information Unit”, they model the problem of retrieving web pages as the minimum-weighted group Steiner tree problem, which is an NP-hard problem. It is rather difficult to identify Steiner trees over large graphs. Also, their method employs a heuristic method to identify the *top-k* answers, and it may fall into the local optimal point and fail to reach the global optimum. In this paper, we propose an effective search method, EASE, to address the problem.

2.1.2 Semi-structured, Structured and Graph Data

Traditional studies of keyword search over semi-structured data always compute LCAs (lowest common ancestors) or its variants [11, 28] of *content nodes* which directly contain

input keywords and take the subtrees rooted at LCAs as answers. However, the problem as to which subtrees are more meaningful for answering the keyword queries remains open. This is because it is not a straightforward task to decide which subtrees with meaningful and complementary elements besides the content nodes should be used as results in order to meaningfully expand their answerability.

Prior works [2] of keyword search over structured data always identify connected trees with minimal cost in the labeled graph as answers. Called *Steiner trees*, they have nodes which are tuples in the database and links which are primary-foreign-key relationships. The trees are identified with the use of an approximation to the Steiner tree problem. However, it is fairly difficult to extract all the Steiner trees in a large graph, which is NP-hard [2]. Moreover, the Steiner tree problem is difficult to adapt for complicated graph databases [10], e.g., complex biological databases, as it only discovers simple tree structures but cannot identify the more meaningful graph structures with rich structural relationships, such as circles. Although Guo et al. [10] have proposed data topology search to improve search effectiveness, their method is constrained by the input of only two keywords.

Traditionally, the inverted index is employed to answer keyword queries. It has been shown to be effective for text and document-based retrieval. However, it is inadequate for supporting keyword queries over structured, semi-structured and graph data because it is fairly difficult to identify the “best” answers that capture rich structural relationships through the inverted index. To address the above-mentioned issues, we propose an effective graph index to improve search performance in this paper.

2.2 r-Radius Steiner Graph Problem

EASE models unstructured data (e.g., text documents), semi-structured data (e.g., XML databases) and structured data (e.g., relational databases) as graphs, where the nodes are respectively documents, elements and tuples, and the edges are respectively hyperlinks*, parent-child relationships (or IDREFS) and primary-foreign-key relationships. The advantage is obvious as EASE addresses the problem of keyword search over graph data. We now formally define the problem of modeling heterogenous data as graphs.

Inspired by the Steiner tree problem [2], we introduce the *Steiner graph problem*. However, graphs with a larger diameter (which is defined as the longest distance between any two nodes in a graph) are not so meaningful and relevant to queries as users are generally frustrated by large and complex graphs. Consequently, we introduce the *r-radius Steiner graph problem*, which is a more interesting and challenging problem of identifying meaningful Steiner graphs with acceptable sizes. To formally describe this problem, we first present several concepts as follows.

DEFINITION 1. (CENTRIC DISTANCE) Given graph \mathcal{G} and any node v in \mathcal{G} , the *centric distance* of v , denoted as $CD(v)$, is the maximal value among the distances between v and any node u in \mathcal{G} , i.e., $CD(v) = \max_{u \in \mathcal{G}} \{D(v, u)\}$, where $D(v, u)$ denotes the distance between v and u , i.e., the length of the shortest path between v and u .

*There may be a large number of hyperlinks among pages on the Web; however, to extract more meaningful and relevant answers, we can only consider the hyperlinks between pages in the same domain.

Table 1: A publication database

Schema		AID		PID		PID	
Authors	Author-Paper	Papers	Paper-Reference				
						citedPID	
<u>Authors</u>		<u>Paper-Reference</u>		<u>Author-Paper</u>			
AID	Name	PID	citedPID	AID	PID		
a1	J. Shanmugasundaram	p1	p2	a1	p2		
a2	L. Guo	p2	p3	a2	p1		
a3	V. Hristidis	p3	p4	a3	p4		
a4	Y. Papakonstantinou	p4	p5	a3	p5		
a5	A. Balmin	p5	p6	a4	p5		
				a4	p6		
				a4	p7		
				a5	p6		
<u>Papers</u>							
PID	Title						
p1	Topology Search over Biological Databases						
p2	XRANK: Ranked Keyword Search over XML Documents						
p3	Bidirectional Expansion for Keyword Search on Graphs						
p4	Finding top-k Answers in Keyword Proximity Search						
p5	Efficient IR-Style Keyword Search over Relational Databases						
p6	Keyword Proximity Search on XML Graphs						
p7	DISCOVER: Keyword Search in Relational Databases						

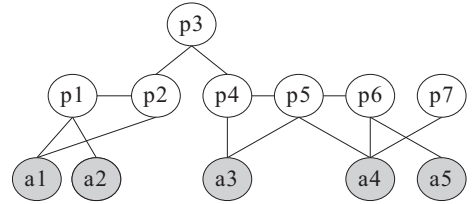


Figure 2: The graph model for the publication database in Table 1

DEFINITION 2. (RADIUS) The *radius* of a graph \mathcal{G} , denoted as $\mathcal{R}(\mathcal{G})$, is the minimal value among the centric distances of every node in \mathcal{G} , i.e., $\mathcal{R}(\mathcal{G}) = \min_{v \in \mathcal{G}} \{CD(v)\}$. \mathcal{G} is called an *r-radius graph* if the radius of \mathcal{G} is exactly r .

DEFINITION 3. (r-RADIUS STEINER GRAPH) Given an *r-radius graph* \mathcal{G} and a keyword query \mathcal{K} . Node w in \mathcal{G} is called a *content node* if w directly contains some input keywords in \mathcal{K} . Node s in \mathcal{G} is called a *Steiner node* if there exist two content nodes, u and v , and s is on the path $u \rightsquigarrow v$ (s may be u or v), where $u \rightsquigarrow v$ denotes a path between u and v . The subgraph of \mathcal{G} composed of the Steiner nodes and associated edges is called an *r-radius Steiner graph*. The radius of an *r-radius Steiner graph* may be smaller than r but cannot be larger than r .

Obviously, we can take the *r-radius graphs* that contain all or a portion of the input keywords as answers, as the *r-radius graphs* are very compact and meaningful, and also contain some relevant and complementary nodes for expanding their answerability. Moreover, *r-radius Steiner graphs* are more concise since non-Steiner nodes are excluded. Although we can take either *r-radius graphs* or *r-radius Steiner graphs* as answers of keyword search over graph data, we adopt the latter option in this paper.

However, it is very difficult to identify all *r-radius Steiner graphs* from a large graph, and hence, we will introduce an effective index later to improve the efficiency of extracting *r-radius Steiner graphs*. Here, to better explain our proposal, we give a running example as described in EXAMPLE 1.

Table 2: Summary of notation

Notation	Descriptions
\mathcal{G}	a graph
$ \mathcal{G} $	the number of nodes in \mathcal{G}
\mathcal{SG}	a Steiner graph
\mathcal{K}	an input keyword query
\mathcal{M}	the adjacency matrix w.r.t. \mathcal{G}
\mathcal{M}^r	the r -th power of \mathcal{M}
\mathcal{N}_i^r (or $\mathcal{N}_{v_i}^r$)	$\{v_j \mathcal{M}_{ij}^r = 1\}$
\mathcal{G}_i^r (or $\mathcal{G}_{v_i}^r$)	the subgraph of \mathcal{G} composed of the nodes in \mathcal{N}_i^r
$v_i \rightsquigarrow v_j$	there is a path between v_i and v_j
$\mathcal{D}(v_i, v_j)$	the distance of v_i and v_j
$\mathcal{CD}(v)$	$\max_{u \in \mathcal{G}} \{\mathcal{D}(v, u)\}$
$\mathcal{R}(\mathcal{G})$	the radius of \mathcal{G} and $\mathcal{R}(\mathcal{G}) = \min_{v \in \mathcal{G}} \{\mathcal{CD}(v)\}$
$v_i \rightsquigarrow^d v_j$	$v_i \rightsquigarrow v_j$ and $\mathcal{D}(v_i, v_j) \leq d$
$\mathcal{G}_i \trianglelefteq \mathcal{G}_j$	\mathcal{G}_i is a subgraph of \mathcal{G}_j (or \mathcal{G}_j is a super-graph of \mathcal{G}_i)
$\mathcal{G}_i \triangleleft \mathcal{G}_j$	$\mathcal{G}_i \trianglelefteq \mathcal{G}_j$ and $\mathcal{G}_i \neq \mathcal{G}_j$
$\mathcal{S}(\mathcal{G}_i, \mathcal{G}_j)$	the similarity between \mathcal{G}_i and \mathcal{G}_j

EXAMPLE 1. Consider the database in Table 1. We model it to a graph \mathcal{G} as illustrated in Figure 2. $\mathcal{D}(a1, a3) = 4$, $\mathcal{CD}(p2) = 5$ and $\mathcal{R}(\mathcal{G}) = 4$. Given a query “IR, Hristidis” on the database in Table 1, we compute the Steiner graph composed of $p4, p5$ and $a3$ with associated edges between them as an answer. This differs from the Steiner tree (i.e., $p5 - a3$) of prior studies, which can cause the loss of meaningful information, especially in databases with complicated structures.

We now formally state the r -radius Steiner graph problem for identifying the most relevant subgraphs with acceptable sizes to answer keyword queries over graph data.

THE r -RADIUS STEINER GRAPH PROBLEM: Given a graph \mathcal{G} and an input keyword query \mathcal{K} , the r -RADIUS STEINER GRAPH PROBLEM is to find all the r -radius Steiner graphs in \mathcal{G} , which contain all or a portion of the input keywords in \mathcal{K} , ranked by relevancy with \mathcal{K} .

As users are usually interested in the *top-k* answers, we mainly discuss how to identify *top-k* r -radius Steiner graphs with the highest scores.

3. EASE: AN ADAPTIVE SEARCH METHOD

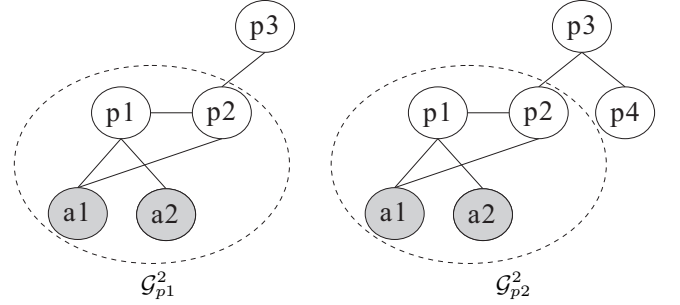
The efficiency and advantages of using inverted indices for facilitating the computation of the “best” answers for online keyword queries are well recognized. However, the inverted indices are not effective for discovering the much richer structural relationships existing in databases with complicated structures [10]. It is therefore important to be able to efficiently and effectively discover these structural relationships, and index them for fast and accurate response. Intuitively, a straightforward way is to enumerate all the combinations of keywords, compute the corresponding r -radius Steiner graphs for each combination, and index these graphs. However, it is prohibitively expensive to discover all these structures since the number of combinations of all keywords in real databases is very large.

Consequently, we propose an effective strategy to discover a portion of the r -radius graphs such that the number of which is proportional to the number of nodes in the graph, and we only need to index and materialize these graphs. More importantly, all of the r -radius graphs can be effectively identified through the indexed ones. In a later section, we will address the issue of extracting r -radius Steiner graphs on the fly by removing non-Steiner nodes from the corresponding indexed r -radius graphs.

Table 3: Adjacency matrix of the graph for the publication databases

		(a) \mathcal{M}											
		a1	p1	a2	p2	p3	p4	a3	p5	a4	p6	p7	a5
a1		1	1	0	1	0	0	0	0	0	0	0	0
p1		1	1	1	1	0	0	0	0	0	0	0	0
a2		0	1	1	0	0	0	0	0	0	0	0	0
p2		1	1	0	1	1	0	0	0	0	0	0	0
p3		0	0	0	1	1	1	0	0	0	0	0	0
p4		0	0	0	0	1	1	1	1	0	0	0	0
a3		0	0	0	0	0	1	1	1	1	0	0	0
p5		0	0	0	0	0	1	1	1	1	1	0	0
a4		0	0	0	0	0	0	0	1	1	1	1	0
p6		0	0	0	0	0	0	0	1	1	1	0	1
p7		0	0	0	0	0	0	0	0	1	0	1	0
a5		0	0	0	0	0	0	0	0	0	1	0	1

		(b) \mathcal{M}^2											
		a1	p1	a2	p2	p3	p4	a3	p5	a4	p6	p7	a5
a1		1	1	1	1	1	0	0	0	0	0	0	0
p1		1	1	1	1	1	0	0	0	0	0	0	0
a2		1	1	1	1	0	0	0	0	0	0	0	0
p2		1	1	1	1	1	1	0	0	0	0	0	0
p3		1	1	0	1	1	1	1	0	0	0	0	0
p4		0	0	0	1	1	1	1	1	1	1	0	0
a3		0	0	0	0	1	1	1	1	1	1	0	0
p5		0	0	0	0	1	1	1	1	1	1	1	1
a4		0	0	0	0	0	1	1	1	1	1	1	1
p6		0	0	0	0	0	1	1	1	1	1	1	1
p7		0	0	0	0	0	0	0	1	1	1	1	0
a5		0	0	0	0	0	0	0	1	1	1	0	1


Figure 3: Two 2-radius graphs

3.1 Adjacency Matrix

In order to efficiently extract r -radius graphs from a given graph $\mathcal{G}(\mathcal{R}(\mathcal{G}) \geq r)^\dagger$, we introduce the concept of *Adjacency Matrix*, $\mathcal{M} = (m_{ij})_{n \times n}$, with respect to \mathcal{G} , which is an $n \times n$ boolean matrix. In \mathcal{M} , the element m_{ij} is 1, iff, there is an edge between v_i and v_j in \mathcal{G} , i.e., $v_i \rightsquigarrow^1 v_j$, where $v_i(v_j)$ denotes the node at the i -th(j -th) row or column in \mathcal{M} while $v_i \rightsquigarrow^d v_j$ denotes that there is a path between v_i and v_j with distance no larger than d ; otherwise, m_{ij} is 0 (m_{ii} is always 1). Iteratively, $m_{ij}^r = 1$, iff, $v_i \rightsquigarrow^r v_j$, where

$\mathcal{M}^r = \overbrace{\mathcal{M} \times \cdots \times \mathcal{M}}^r = (m_{ij}^r)_{n \times n}$. To ease the discussion that follows, we summarize the notation we use in Table 2.

\mathcal{M}^r is said to be the r -th power of \mathcal{M} . $\mathcal{N}_i^r = \{v_j | \mathcal{M}_{ij}^r = 1\}$ is the set of nodes which have a path to v_i with distance no larger than r . \mathcal{G}_i^r denotes the subgraph of \mathcal{G} with respect to the i -th row of \mathcal{M}^r , which is composed of the nodes in \mathcal{N}_i^r and the associated edges. $\mathcal{G}_{v_i}^r(\mathcal{N}_{v_i}^r)$ can be interchangeably employed instead of $\mathcal{G}_i^r(\mathcal{N}_i^r)$ if there is no ambiguity. We use $\mathcal{G}_i \trianglelefteq \mathcal{G}_j$ to denote that \mathcal{G}_i is a subgraph of \mathcal{G}_j (equivalently,

[†]Without loss of generality, we suppose \mathcal{G} is a connected graph, the radius of which is no smaller than r , i.e., $\mathcal{R}(\mathcal{G}) \geq r$. Even if \mathcal{G} is an unconnected graph, we can decompose it into a set of connected graphs and thus EASE can adapt to the unconnected graphs.

\mathcal{G}_j is also called a super-graph of \mathcal{G}_i). $\mathcal{G}_i \triangleleft \mathcal{G}_j$ denotes that \mathcal{G}_i is a proper subgraph of \mathcal{G}_j , i.e., $\mathcal{G}_i \triangleleft \mathcal{G}_j$ and $\mathcal{G}_i \neq \mathcal{G}_j$. $|\mathcal{G}|$ denotes the number of nodes in \mathcal{G} .

EXAMPLE 2. We can construct the adjacency matrix of the graph in Figure 2 as illustrated in Table 3. We note that $v_1=a_1$, $v_2=p_1$, $v_3=a_2$ and $v_4=p_2$. $\mathcal{N}_2^2=\mathcal{N}_{p_1}^2=\{a_1,p_1,a_2,p_2,p_3\}$ and $\mathcal{N}_4^2=\mathcal{N}_{p_2}^2=\{a_1,p_1,a_2,p_2,p_3,p_4\}$. The two sub-graphs $\mathcal{G}_2^2(\mathcal{G}_{p_1}^2)$ and $\mathcal{G}_4^2(\mathcal{G}_{p_2}^2)$ are illustrated in Figure 3.

To effectively extract r -radius graphs according to the adjacency matrix, we provide LEMMA 1 to help determine the subgraphs in \mathcal{G} that are r -radius graphs.

LEMMA 1. Given a graph $\mathcal{G}(\mathcal{R}(\mathcal{G}) \geq r > 1)$, $\forall i, 1 \leq i \leq |\mathcal{G}|$, $\mathcal{G}_{v_i}^r$ is an r -radius graph, if, $\forall v_k \in \mathcal{N}_{v_i}^r$, $\mathcal{N}_{v_i}^r \not\subseteq \mathcal{N}_{v_k}^{r-1}$.

PROOF. As all the nodes in $\mathcal{N}_{v_i}^r$ connect to v_i and all the nodes with distances no larger than r to v_i are in $\mathcal{N}_{v_i}^r$, the nodes in $\mathcal{N}_{v_i}^r$ and the edges associated with them can construct a subgraph of \mathcal{G} , $\mathcal{G}_{v_i}^r$. We need to prove $\mathcal{R}(\mathcal{G}_{v_i}^r)=r$.

We first prove that $\mathcal{R}(\mathcal{G}_{v_i}^r) \leq r$.

$\forall u \in \mathcal{N}_{v_i}^r$, $\mathcal{D}(v_i, u) \leq r$ according to the definition of \mathcal{M}^r , and thus $\mathcal{CD}(v_i) = \max_{u \in \mathcal{G}_{v_i}^r} \{\mathcal{D}(v_i, u)\} \leq r$. Hence, based on DEFINITION 2, we have $\mathcal{R}(\mathcal{G}_{v_i}^r) = \min_{v \in \mathcal{G}_{v_i}^r} \{\mathcal{CD}(v)\} \leq \mathcal{CD}(v_i) \leq r$.

We then prove that $\mathcal{R}(\mathcal{G}_{v_i}^r) \geq r$.

$\forall v_k \in \mathcal{N}_{v_i}^r$, as $\mathcal{N}_{v_i}^r \not\subseteq \mathcal{N}_{v_k}^{r-1}$, there must exist a node u_k , $u_k \in \mathcal{N}_{v_i}^r$ and $u_k \notin \mathcal{N}_{v_k}^{r-1}$, thus $\mathcal{D}(v_k, u_k) \geq r$ (otherwise, if $\mathcal{D}(v_k, u_k) < r$, $u_k \in \mathcal{N}_{v_k}^{r-1}$, which contradicts $u_k \notin \mathcal{N}_{v_k}^{r-1}$). Thus, $\forall v_k \in \mathcal{N}_{v_i}^r$, we have $\mathcal{CD}(v_k) = \max_{u \in \mathcal{G}_{v_i}^r} \{\mathcal{D}(v_k, u)\} \geq \mathcal{D}(v_k, u_k) \geq r$ in $\mathcal{G}_{v_i}^r$. Accordingly, $\mathcal{R}(\mathcal{G}_{v_i}^r) = \min_{v_k \in \mathcal{G}_{v_i}^r} \{\mathcal{CD}(v_k)\} \geq r$.

Therefore, $\mathcal{R}(\mathcal{G}_{v_i}^r)=r$ and $\mathcal{G}_{v_i}^r$ is an r -radius graph. \square

As formalized in LEMMA 1, we can determine whether the subgraph $\mathcal{G}_{v_i}^r$ with respect to the i -th row of \mathcal{M}^r is an r -radius graph. For example, $\mathcal{G}_{p_2}^2$ is a 2-radius graph while $\mathcal{G}_{a_2}^2$ is not as $\mathcal{N}_{a_2}^2 \subseteq \mathcal{N}_{p_1}^1$ as shown in Table 3.

In order to extract all of the r -radius Steiner graphs from \mathcal{G} , we must prove that any r -radius Steiner graph in \mathcal{G} corresponds to a subgraph $\mathcal{G}_{v_i}^r$, and we provide THEOREM 1 for extracting r -radius Steiner graphs.

THEOREM 1. Suppose $\{\mathcal{SG}_1^r, \mathcal{SG}_2^r, \dots, \mathcal{SG}_p^r\}$ is the set of the r -radius Steiner graphs with respect to a graph $\mathcal{G}(\mathcal{R}(\mathcal{G}) \geq r)$ and a keyword query, $\forall 1 \leq j \leq p$, $\exists 1 \leq i \leq |\mathcal{G}|$, $\mathcal{SG}_j^r \triangleleft \mathcal{G}_{v_i}^r$ holds.

PROOF. $\forall \mathcal{SG}_j^r$, there must exist a corresponding r -radius graph \mathcal{G}_{r_j} according to DEFINITION 3, such that $\mathcal{SG}_j^r \triangleleft \mathcal{G}_{r_j}$. We then prove that $\exists \mathcal{G}_{v_i}^r$, $\mathcal{G}_{r_j} \triangleleft \mathcal{G}_{v_i}^r$.

As \mathcal{G}_{r_j} is an r -radius graph, there must exist a node $v_i \in \mathcal{G}_{r_j}$ and $\mathcal{CD}(v_i)=r$. Since for any node $u \in \mathcal{G}_{r_j}$, $\mathcal{D}(v_i, u) \leq r$, thus u must be in the node set of $\mathcal{G}_{v_i}^r$ (i.e., $\mathcal{N}_{v_i}^r$) according to the definitions of $\mathcal{G}_{v_i}^r$ and \mathcal{M}^r . Hence, all the nodes in \mathcal{G}_{r_j} must be in $\mathcal{G}_{v_i}^r$. Thus, $\mathcal{G}_{r_j} \triangleleft \mathcal{G}_{v_i}^r$.

Hence, $\mathcal{SG}_j^r \triangleleft \mathcal{G}_{r_j} \triangleleft \mathcal{G}_{v_i}^r$. \square

Based on THEOREM 1, we can extract r -radius Steiner graphs through the adjacency matrix. To facilitate efficient retrieval of r -radius graphs, we construct a novel *graph index*. The entries of the graph index are terms contained in the graph and each entry preserves the r -radius graphs that contain the term. To construct the graph index, we first extract r -radius graphs as stated in LEMMA 1, then for each term k_i , we keep the set of all r -radius graphs that contain k_i , denoted as \mathcal{I}_{k_i} , i.e., $\mathcal{I}_{k_i} = \{\mathcal{G}_{v_j}^r | \mathcal{G}_{v_j}^r \text{ contains } k_i\}$.

To process a keyword query $\mathcal{K}=\{k_1, k_2, \dots, k_m\}$, we first retrieve the set \mathcal{I}_{k_i} of those r -radius graphs which contain k_i based on the graph index, and then union every \mathcal{I}_{k_i} to compute $\cup_{i=1}^m \mathcal{I}_{k_i}$, which is the set of r -radius graphs that contain all or a portion of the keywords in \mathcal{K}^\ddagger . Finally, we extract the r -radius Steiner graphs by removing the non-Steiner nodes from the corresponding r -radius graphs, and rank the results to return the *top-k* answers.

Given an r -radius graph \mathcal{G}^r and its *content nodes*, c_1, c_2, \dots, c_q , which directly contain some of the input keywords, we compute the Steiner nodes and construct the corresponding r -radius Steiner graph as follows:

- i) Compute $\mathcal{P}(c_i)$, the set of those nodes which have a path to c_i in \mathcal{G}_i , i.e., $\mathcal{P}(c_i) = \{u | u \rightsquigarrow c_i \text{ in } \mathcal{G}_i\}$, where \mathcal{G}_i is the subgraph of \mathcal{G}^r , by removing the nodes in $\{c_1, c_2, \dots, c_{i-1}, c_{i+1}, \dots, c_q\}$ and the associated edges.
- ii) Extract the set of Steiner nodes in \mathcal{G}^r , \mathcal{P} , where $\mathcal{P} = \cup_{i=1}^q \cup_{j=i+1}^q (\mathcal{P}(c_i) \cap \mathcal{P}(c_j)) \cup \{c_1, c_2, \dots, c_q\}$.
- iii) Construct the r -radius Steiner graph, i.e., the subgraph of \mathcal{G}^r , which is composed of the nodes in \mathcal{P} and the associated edges.

EXAMPLE 3. Consider the database described in Table 1, we obtain its adjacency matrix \mathcal{M} and compute \mathcal{M}^r as shown in Table 3 (r is set to 2 in the remaining examples throughout this paper). Subsequently, we derive the r -radius graphs, e.g., $\mathcal{G}_{p_1}^r$ and $\mathcal{G}_{p_2}^r$ based on the second and fourth rows of \mathcal{M}^r respectively, which are illustrated in Figure 3. We note that $\mathcal{G}_{p_1}^r \triangleleft \mathcal{G}_{p_2}^r$. To answer the keyword query "Shanmugasundaram, Guo, XRANK", we first retrieve the r -radius graphs, $\mathcal{G}_{p_1}^r$ and $\mathcal{G}_{p_2}^r$, which contain the three keywords, based on the graph index, and then extract the corresponding r -radius Steiner graphs, i.e., the circled subgraphs as illustrated in Figure 3, by removing the non-Steiner nodes.

From EXAMPLE 3, we observe that some r -radius graphs are contained in others. As a further illustration, we get the following expressions from the graph in Figure 2: i) $\mathcal{G}_{a_2}^r \triangleleft \mathcal{G}_{a_1}^r = \mathcal{G}_{p_1}^r \triangleleft \mathcal{G}_{p_2}^r$; ii) $\mathcal{G}_{a_3}^r \triangleleft \mathcal{G}_{p_4}^r$; iii) $\mathcal{G}_{a_5}^r \triangleleft \mathcal{G}_{p_6}^r = \mathcal{G}_{a_4}^r \triangleleft \mathcal{G}_{p_5}^r$; and iv) $\mathcal{G}_{p_7}^r \triangleleft \mathcal{G}_{p_6}^r = \mathcal{G}_{a_4}^r \triangleleft \mathcal{G}_{p_5}^r$. Consequently, it is sufficient for us just to keep the graphs $\mathcal{G}_{p_2}^r, \mathcal{G}_{p_3}^r, \mathcal{G}_{p_4}^r$ and $\mathcal{G}_{p_5}^r$ in the graph index. We shall address this problem next.

3.2 Maximal r-Radius Graph

To avoid maintaining redundant overlapping r -radius graphs in the graph index, we introduce the concept of *maximal r-radius graph* in this section.

DEFINITION 4. (MAXIMAL r -RADIUS GRAPH) Given a graph \mathcal{G} and an r -radius subgraph $\mathcal{G}_{v_i}^r$ in \mathcal{G} , $\mathcal{G}_{v_i}^r$ is called a *maximal r-radius subgraph* if there is no other r -radius subgraph that contains $\mathcal{G}_{v_i}^r$.

Based on DEFINITION 4, we only need to keep the maximal r -radius graphs in the graph index to minimize storage without affecting the search results. This is because all other r -radius graphs can be reconstructed from their corresponding super-graphs. For example, $\mathcal{G}_{p_2}^r, \mathcal{G}_{p_3}^r, \mathcal{G}_{p_4}^r$ and $\mathcal{G}_{p_5}^r$ are maximal r -radius graphs and will be kept in the graph index while other graphs, e.g., $\mathcal{G}_{p_1}^r$ and $\mathcal{G}_{a_4}^r$, can be reconstructed when necessary.

In fact, the maximal r -radius graphs can be directly extracted from \mathcal{M}^r , as captured by LEMMA 2 below.

\ddagger If we consider "AND" predicate, we merge \mathcal{I}_{k_i} to compute $\cap_{i=1}^m \mathcal{I}_{k_i}$, which is the set of r -radius graphs that contain all keywords.

Table 4: Graph similarities

	\mathcal{G}_{p2}^r	\mathcal{G}_{p3}^r	\mathcal{G}_{p4}^r	\mathcal{G}_{p5}^r
\mathcal{G}_{p2}^r	1	5/8	3/10	1/6
\mathcal{G}_{p3}^r	-	1	5/9	4/11
\mathcal{G}_{p4}^r	-	-	1	2/3
\mathcal{G}_{p5}^r	-	-	-	1

LEMMA 2. Given a graph $\mathcal{G}(\mathcal{R}(\mathcal{G}) \geq r)$, $\forall i, 1 \leq i \leq |\mathcal{G}|$, $\mathcal{G}_{v_i}^r$ is a maximal r -radius graph if $\forall k \in \{t | \mathcal{M}_{it}^r = 1, \mathcal{N}_i^r \neq \mathcal{N}_k^r\}$, $\mathcal{N}_i^r \not\subset \mathcal{N}_k^r$.

PROOF. We first prove that $\mathcal{G}_{v_i}^r$ is an r -radius graph.

i) $\forall k \in \{t | \mathcal{M}_{it}^r = 1, \mathcal{N}_i^r \neq \mathcal{N}_k^r\}$, as $\mathcal{N}_i^r \not\subset \mathcal{N}_k^r$, there must exist $u_k, u_k \in \mathcal{N}_i^r$ and $u_k \notin \mathcal{N}_k^r$, and thus we have $\mathcal{D}(v_k, u_k) > r$ (otherwise, if $\mathcal{D}(v_k, u_k) \leq r$, u_k must be in \mathcal{N}_k^r , which contradicts $u_k \notin \mathcal{N}_k^r$). Hence, $\mathcal{CD}(v_k) = \max_{u \in \mathcal{G}_{v_i}^r} \{\mathcal{D}(v_k, u)\} \geq \mathcal{D}(v_k, u_k) > r$.

ii) $\forall j \in \{t | \mathcal{M}_{it}^r = 1, \mathcal{N}_i^r = \mathcal{N}_j^r\}$, as $\mathcal{R}(\mathcal{G}) \geq r$, $\mathcal{CD}(v_j) = r$ in $\mathcal{G}_{v_i}^r$.

Thus, $\forall v_k \in \mathcal{G}_{v_i}^r$, $\mathcal{CD}(v_k) \geq r$ and $\mathcal{CD}(v_i) = r$. Hence, we have $\mathcal{R}(\mathcal{G}_{v_i}^r) = \min_{v_k \in \mathcal{G}_{v_i}^r} \{\mathcal{CD}(v_k)\} = r$, and $\mathcal{G}_{v_i}^r$ is an r -radius graph.

We then prove that $\mathcal{G}_{v_i}^r$ must be a maximal r -radius graph.

As $\forall k \in \{t | \mathcal{M}_{it}^r = 1, \mathcal{N}_i^r \neq \mathcal{N}_k^r\}$, $\mathcal{N}_i^r \not\subset \mathcal{N}_k^r$, there cannot exist another r -radius graph that is a super-graph of $\mathcal{G}_{v_i}^r$. Thus, $\mathcal{G}_{v_i}^r$ must be a maximal r -radius graph. \square

There are still some overlaps between different maximal r -radius graphs. For example, \mathcal{G}_{p2}^r and \mathcal{G}_{p3}^r both contain the nodes, a_1, p_1, p_2, p_3 and p_4 . To further reduce storage and improve search performance, we introduce a technique for clustering the maximal r -radius graphs in Section 3.3.

3.3 Graph Partitioning

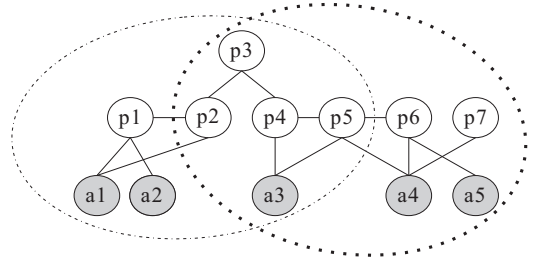
Conventional Steiner tree based methods [2] typically need to traverse the graph initializing from the content nodes to identify the Steiner trees with the minimal cost. They have to maintain the whole graph in memory, which is not practical for large graphs. To alleviate the problem, we propose a graph partitioning based method to achieve the needed search efficiency.

We cluster r -radius graphs so that we can partition the graph to facilitate identifying r -radius graphs. We first cluster the r -radius graphs and then partition the whole graph based on the clusters. Each cluster corresponds to a portion of the graph. Clustering maximal r -radius graphs has the following salient advantages: (i) We only need to maintain a physical graph for each cluster while all the maximal r -radius graphs only preserve their nodes instead of the corresponding overlapping graphs. This avoids the incurrence of huge storage, and is similar to the views on top of underlying physical tables in RDBMS, in which clusters correspond to physical tables while maximal r -radius graphs correspond to views. (ii) We only need to retrieve the corresponding relevant graph partitions instead of maintaining the whole graph in order to identify the r -radius graphs.

To meaningfully cluster r -radius graphs, we first define the similarity between any two graphs and then cluster the maximal r -radius graphs based on their similarities.

DEFINITION 5. (GRAPH SIMILARITY) Given two maximal r -radius subgraphs \mathcal{G}_i^r and \mathcal{G}_j^r in a given graph, their graph similarity, $\mathcal{S}(\mathcal{G}_i^r, \mathcal{G}_j^r)$, is $\frac{|\mathcal{N}_i^r \cap \mathcal{N}_j^r|}{|\mathcal{N}_i^r \cup \mathcal{N}_j^r|}$, where \mathcal{N}_i^r and \mathcal{N}_j^r denote the node sets of \mathcal{G}_i^r and \mathcal{G}_j^r respectively.

A bigger overlap between the nodes of the two graphs implies a larger graph similarity between them, and consequently, a higher probability that they will be clustered



$$\mathcal{G}_{p2}^r + \mathcal{G}_{p3}^r \text{ (cluster1)} \quad \mathcal{G}_{p4}^r + \mathcal{G}_{p5}^r \text{ (cluster2)}$$

Figure 4: Two clusters

together. It is obvious that the graph similarity scales well with the number of overlapping nodes. Moreover, given two maximal r -radius subgraphs, employing the number of their overlapping nodes is sound as the edges associated with the overlapping nodes are also the same. In addition, we note that the graph similarity preserves the following properties:

- SYMMETRY. $\mathcal{S}(\mathcal{G}_i, \mathcal{G}_j) = \mathcal{S}(\mathcal{G}_j, \mathcal{G}_i)$;
- POSITIVITY. $0 \leq \mathcal{S}(\mathcal{G}_i, \mathcal{G}_j) \leq 1$, for any \mathcal{G}_i and \mathcal{G}_j ;
- REFLEXIVITY. $\mathcal{S}(\mathcal{G}_i, \mathcal{G}_j) = 1$ iff $\mathcal{G}_i = \mathcal{G}_j$;

which indicate that graph similarity is a good metric to evaluate the similarity between any two graphs. Based on the graph similarity, we can now cluster the maximal r -radius subgraphs by employing an existing method such as the K-mean, K-medoids and EM algorithms.

To effectively compute the graph similarity between two maximal r -radius graphs, we introduce LEMMA 3.

LEMMA 3. Given two maximal r -radius subgraphs w.r.t. \mathcal{M}^r , \mathcal{G}_i^r and \mathcal{G}_j^r , $\mathcal{S}(\mathcal{G}_i^r, \mathcal{G}_j^r) = \frac{|\{k | \mathcal{M}_{ik}^r = 1 \text{ and } \mathcal{M}_{jk}^r = 1\}|}{|\{k | \mathcal{M}_{ik}^r = 1 \text{ or } \mathcal{M}_{jk}^r = 1\}|}$.

Following LEMMA 3, we can compute the graph similarity based on the adjacency matrix, which is much easier to manipulate than the original graphs.

EXAMPLE 4. Consider the database in Table 1, we note that the four graphs $\mathcal{G}_{p2}^r, \mathcal{G}_{p3}^r, \mathcal{G}_{p4}^r$ and \mathcal{G}_{p5}^r are the maximal r -radius graphs. We can compute their graph similarities as shown in Table 4 based on the adjacency matrix in Table 3. If we cluster the four maximal r -radius graphs into two clusters, \mathcal{G}_{p2}^r and \mathcal{G}_{p3}^r will fall into the same cluster while \mathcal{G}_{p4}^r and \mathcal{G}_{p5}^r will be in another cluster as illustrated in Figure 4. On the other hand, if we cluster them into three clusters, \mathcal{G}_{p4}^r and \mathcal{G}_{p5}^r will fall into the same cluster while \mathcal{G}_{p2}^r and \mathcal{G}_{p3}^r are in the other two clusters respectively.

To summarize, given a graph, we first obtain its adjacency matrix \mathcal{M} and compute \mathcal{M}^r . We then extract the maximal r -radius graphs according to LEMMA 2 and compute the graph similarities between any two maximal r -radius graphs based on LEMMA 3. Subsequently, we cluster the graphs by employing the existing K-means algorithm and partition the graph. Finally, we construct the graph index to materialize the maximal r -radius graphs. To illustrate, we consider the example below, which computes the set of r -radius Steiner graphs that contain all or a portion of the input keywords.

EXAMPLE 5. Given the database in Table 1 and a keyword query "DISCOVER, Relational, Database, Papakonstantinou", we first retrieve the keyword lists based on the graph index (Table 5). We then derive the set of $\{\mathcal{G}_{p2}^r, \mathcal{G}_{p3}^r, \mathcal{G}_{p4}^r, \mathcal{G}_{p5}^r\}$,

Table 5: Graph index

Terms	r -radius graphs : \mathcal{I}_{k_i}
Database	$\mathcal{G}_{p_2}^r, \mathcal{G}_{p_3}^r, \mathcal{G}_{p_4}^r, \mathcal{G}_{p_5}^r$
DISCOVER	$\mathcal{G}_{p_5}^r$
Papakonstantinou	$\mathcal{G}_{p_4}^r, \mathcal{G}_{p_5}^r$
Relational	$\mathcal{G}_{p_3}^r, \mathcal{G}_{p_4}^r, \mathcal{G}_{p_5}^r$

where each graph contains some input keywords. $\mathcal{G}_{p_5}^r$ contains all of the input keywords. Finally, we refine the maximal r -radius graphs on top of the corresponding graph partitions (instead of traversing the whole graph) to obtain the r -radius Steiner graphs. For instance, we can get $\mathcal{S}\mathcal{G}_{p_5}^r$ by removing the non-Steiner nodes, e.g., p_3, p_4, a_3 and a_5 , from $\mathcal{G}_{p_5}^r$, as shown in Figure 5.

4. RANKING FUNCTIONS

In this section, we first discuss how to meaningfully rank r -radius Steiner graphs and identify the *top-k* answers based on the existing proposals. Next, we propose a new measure based on the structural compactness between content nodes and the structural relevancy between input keywords with respect to an r -radius Steiner graph.

4.1 TF • IDF-based IR Ranking

The basic idea of the ranking method used in the existing literature, such as [15, 20, 22], is to first assign each r -radius graph a score using a standard IR-ranking formula (or its variants), and then combine the individual scores using a score aggregation function, such as SUM, to obtain the final score. For example, the TF-IDF-based IR-style ranking function weights an r -radius Steiner graph by considering textual relevancy in IR literature, which takes into account term frequency (**tf**), inverse document frequency (**idf**) and normalized document length (**ndl**). **tf** and **idf** are well employed to rank documents in the IR literature while **ndl** is used to normalize document length as a longer document has a higher likelihood to contain many more keywords. We can compute the three parameters as follows:

$$ntf_{(k_i, \mathcal{G})} = 1 + \ln(1 + \ln(1 + tf_{(k_i, \mathcal{G})})) \quad (1)$$

$$idf_{k_i} = \ln \frac{N + 1}{N_{k_i} + 1} \quad (2)$$

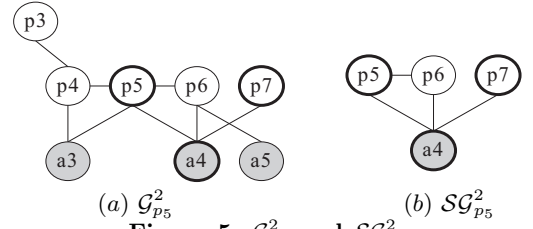
$$ndl_{\mathcal{G}} = (1 - s) + s * \frac{tl_{\mathcal{G}}}{avg_{tl}} \quad (3)$$

where $tf_{(k_i, \mathcal{G})}$ in Equation 1 denotes the term frequency of k_i in \mathcal{G} ; N and N_{k_i} in Equation 2 respectively denote the number of maximal r -radius graphs and the number of those maximal r -radius graphs that contain k_i . $tl_{\mathcal{G}}$ in Equation 3 denotes the total number of terms in \mathcal{G} and avg_{tl} is the average number of terms among all such r -radius graphs while s is a parameter taken from IR literature, which has been extensively discussed and typically set to 0.2 [20]. We combine the three parameters to evaluate the document relevancy between an input keyword k_i and a given Steiner graph $\mathcal{S}\mathcal{G}$, denoted as $\text{SCORE}_{\text{IR}}(k_i, \mathcal{S}\mathcal{G})$ as formalized in Equation 4, where \mathcal{G} is the corresponding r -radius graph w.r.t. $\mathcal{S}\mathcal{G}$ [§].

$$\text{SCORE}_{\text{IR}}(k_i, \mathcal{S}\mathcal{G}) = \frac{ntf_{(k_i, \mathcal{G})} * idf_{k_i}}{ndl_{\mathcal{G}}} \quad (4)$$

Based on the textual relevancy of k_i in $\mathcal{S}\mathcal{G}$, we compute the overall score between an input keyword query \mathcal{K} and $\mathcal{S}\mathcal{G}$ by summing up $\text{SCORE}_{\text{IR}}(k_i, \mathcal{S}\mathcal{G})$, as shown in Equation 5.

[§] $ntf_{(k_i, \mathcal{S}\mathcal{G})} = ntf_{(k_i, \mathcal{G})}$; $tf_{(k_i, \mathcal{S}\mathcal{G})} = tf_{(k_i, \mathcal{G})}$ for each input keyword k_i .


Figure 5: $\mathcal{G}_{p_5}^2$ and $\mathcal{S}\mathcal{G}_{p_5}^2$

$$\text{SCORE}_{\text{IR}}(\mathcal{K}, \mathcal{S}\mathcal{G}) = \sum_{k_i \in \mathcal{K}} \text{SCORE}_{\text{IR}}(k_i, \mathcal{S}\mathcal{G}) \quad (5)$$

Although the TF-IDF-based ranking methods are efficient for textual documents, they are inefficient for semi-structured and structured data. From the IR perspective, traditional textual relevancy is important. However, due to our use of graph in modeling, the ranking of graph data becomes equally if not more important, and the structural compactness of r -radius Steiner graphs is the essence of the comparison. This is so because identifying rich structural relationships should be at least as important as discovering more keywords, and in some cases, even more crucial. Therefore, we propose a novel ranking function by incorporating structural compactness from the DB point of view.

4.2 Structural Compactness-based DB Ranking

Intuitively, when an r -radius Steiner graph $\mathcal{S}\mathcal{G}$ is more compact, $\mathcal{S}\mathcal{G}$ is more likely to be meaningful and relevant. Accordingly, the structural compactness score should be larger. As such, the compactness of $\mathcal{S}\mathcal{G}$ should include the following parameters: *i*) the structural compactness between content nodes in $\mathcal{S}\mathcal{G}$, and *ii*) the structural relevancy between input keywords w.r.t. $\mathcal{S}\mathcal{G}$. We note that when the length of a path between two content nodes is larger, the relevancy between them is smaller. Further, there may be multiple paths between two content nodes, and we should consider all of them. Based on the above rationale, we propose Equation 6 to score the overall structural compactness between any two content nodes:

$$\text{SIM}(n_i, n_j) = \sum_{n_i \rightsquigarrow n_j} \frac{1}{(|n_i \rightsquigarrow n_j| + 1)^2} \quad (6)$$

where n_i and n_j are two content nodes, $n_i \rightsquigarrow n_j$ is any path between n_i and n_j , and $|n_i \rightsquigarrow n_j|$ is the length of $n_i \rightsquigarrow n_j$. An important feature of $\text{SIM}(n_i, n_j)$ is that it can be pre-computed and materialized off-line, based on the fact that $\text{SIM}_{\mathcal{G}}(n_i, n_j) = \text{SIM}_{\mathcal{S}\mathcal{G}}(n_i, n_j)$ holds as formalized in LEMMA 4 below, where \mathcal{G} denotes the corresponding r -radius graph of $\mathcal{S}\mathcal{G}$ while $\text{SIM}_{\mathcal{G}}(n_i, n_j)$ and $\text{SIM}_{\mathcal{S}\mathcal{G}}(n_i, n_j)$ denote the structural compactness between n_i and n_j in \mathcal{G} and $\mathcal{S}\mathcal{G}$ respectively. It is clear that we can compute the overall structural compactness by summing up several materialized scores. For example, in Figure 5, $\text{SIM}_{\mathcal{S}\mathcal{G}_{p_5}^2}(a_4, p_5) = \text{SIM}_{\mathcal{G}_{p_5}^2}(a_4, p_5) = \frac{1}{(|a_4 \rightsquigarrow p_5| + 1)^2} + \frac{1}{(|a_4 \rightsquigarrow p_6 \rightsquigarrow p_5| + 1)^2} = \frac{1}{4} + \frac{1}{9} = \frac{13}{36}$.

LEMMA 4. Given an r -radius graph \mathcal{G} and its corresponding Steiner graph $\mathcal{S}\mathcal{G}$ with respect to a given keyword query, the following equation holds:

$$\text{SIM}_{\mathcal{G}}(n_i, n_j) = \text{SIM}_{\mathcal{S}\mathcal{G}}(n_i, n_j)$$

where n_i and n_j are any two content nodes in $\mathcal{S}\mathcal{G}$.

PROOF. Note that if there is a path from n_i to n_j in \mathcal{G} , there must be a similar path in $\mathcal{S}\mathcal{G}$ and vice versa based on DEFINITION 3. Thus, $\text{SIM}_{\mathcal{G}}(n_i, n_j) = \text{SIM}_{\mathcal{S}\mathcal{G}}(n_i, n_j)$. \square

Although the structural compactness between two content nodes can measure the structural relevancy of r -radius graphs, it cannot evaluate the structural relevancy among input keywords, which captures the phrase-based relevancy between input keywords. It follows that a smaller distance between input keywords indicates a higher structural relevancy between them. This is particularly so for keywords in the same node that will represent a phrase. We therefore propose Equation 7 to capture this parameter.

$$\text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG}) = \frac{1}{|\mathcal{C}_{k_i} \cup \mathcal{C}_{k_j}|} \sum_{n_i \in \mathcal{C}_{k_i}; n_j \in \mathcal{C}_{k_j}} \text{SIM}(n_i, n_j) \quad (7)$$

where \mathcal{C}_{k_i} denotes the set of all the content nodes that contain k_i in \mathcal{SG} , and $|\mathcal{C}_{k_i}|$ denotes the number of nodes in \mathcal{C}_{k_i} , which is used to normalize the structural relevancy between two input keywords. Consequently, a larger overall structural compactness score of \mathcal{SG} indicates that \mathcal{SG} is more likely to be relevant and meaningful to \mathcal{K} .

We note that the structural relevancy between input keywords has a salient feature that if $\text{SIM}(\langle k_i, k_t \rangle | \mathcal{SG})$ and $\text{SIM}(\langle k_t, k_j \rangle | \mathcal{SG})$ are large, $\text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG})$ must be also large and thus k_i, k_j and k_t must be very relevant to each other w.r.t. \mathcal{SG} . Thus, a key feature of structural relevancy is that we can use the structural relevancy between any two input keywords to capture the relevancy between all of the input keywords as illustrated in Equation 8. This feature can help capture the rich structural information of \mathcal{SG} while the textual relevancy in Equation 5 cannot. Formally, given a keyword query $\mathcal{K}=\{k_1, k_2, \dots, k_m\}$ and an r -radius Steiner graph \mathcal{SG} , the overall structural compactness of \mathcal{SG} w.r.t. \mathcal{K} , denoted as $\text{SCORE}_{\text{DB}}(\mathcal{K}, \mathcal{SG})$, can be computed as follows:

$$\text{SCORE}_{\text{DB}}(\mathcal{K}, \mathcal{SG}) = \sum_{1 \leq i < j \leq m} \text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG}) \quad (8)$$

By taking into account both document relevancy from the IR perspective and structural compactness/relevancy from the DB perspective to capture structural relationships, we present a more accurate function for scoring r -radius Steiner graphs as given in Equation 9.

$$\text{SCORE}(\mathcal{K}, \mathcal{SG}) = \sum_{1 \leq i < j \leq m} \text{SCORE}(\langle k_i, k_j \rangle | \mathcal{SG}) \quad (9)$$

where

$$\text{SCORE}(\langle k_i, k_j \rangle | \mathcal{SG}) = \text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG}) * (\text{SCORE}_{\text{IR}}(k_i, \mathcal{SG}) + \text{SCORE}_{\text{IR}}(k_j, \mathcal{SG})) \quad (10)$$

$\text{SCORE}(\langle k_i, k_j \rangle | \mathcal{SG})$ measures the overall relevancy score of $\langle k_i, k_j \rangle$ in \mathcal{SG} based on the structural compactness/relevancy and IR scores. Note that, $\text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG})$ is taken as the weight of the sum of two IR scores, i.e., $\text{SCORE}_{\text{IR}}(k_i, \mathcal{SG})$ and $\text{SCORE}_{\text{IR}}(k_j, \mathcal{SG})$. A larger $\text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG})$ means that k_i and k_j are more relevant w.r.t. \mathcal{SG} , and thus, the overall score of $\langle k_i, k_j \rangle$ in \mathcal{SG} is expected to be larger.

5. INDEXING

To efficiently identify the $top-k$ answers with the highest scores, we examine the issues of indexing in this section.

Given any two keywords k_i and k_j in the graph, and an r -radius graph \mathcal{SG} , the scores of $\text{SCORE}_{\text{IR}}(k_i, \mathcal{SG})$ and $\text{SCORE}_{\text{IR}}(k_j, \mathcal{SG})$ in Equation 4 and $\text{SIM}(\langle k_i, k_j \rangle | \mathcal{SG})$ in Equation 7 share the key feature that they can be pre-computed and materialized off-line. Based on this observation, we can materialize $\text{SCORE}(\langle k_i, k_j \rangle | \mathcal{SG})$.

Table 6: EI-Index: an extended inverted index

Keyword Pair	$\langle r$ -radius Graph, Score \rangle
$\langle \text{Database, DISCOVER} \rangle$	$\mathcal{G}_{p_5}^r, 1.53$
$\langle \text{Database, Papakonstantinou} \rangle$	$\mathcal{G}_{p_5}^r, 0.38; \mathcal{G}_{p_4}^r, 0.19$
$\langle \text{Database, Relational} \rangle$	$\mathcal{G}_{p_5}^r, 0.85; \mathcal{G}_{p_3}^r, 0.35; \mathcal{G}_{p_4}^r, 0.34$
$\langle \text{DISCOVER, Papakonstantinou} \rangle$	$\mathcal{G}_{p_5}^r, 0.54$
$\langle \text{DISCOVER, Relational} \rangle$	$\mathcal{G}_{p_5}^r, 1.95$
$\langle \text{Papakonstantinou, Relational} \rangle$	$\mathcal{G}_{p_5}^r, 0.57; \mathcal{G}_{p_4}^r, 0.28$
$\langle \dots, \dots \rangle$	\dots

We devise an extended inverted index (EI-Index) to maintain such scores. Different from the traditional inverted index, the entries of EI-Index are keyword pairs (combinations of two keywords), and the values of each entry is the maximal r -radius graphs that contain the keyword pair and the corresponding scores. For example, we can construct the EI-Index of the graph in Figure 2 as illustrated in Table 6.

To answer a keyword query $\mathcal{K}=\{k_1, k_2, \dots, k_m\}$, we first retrieve the maximal r -radius graphs for each keyword pair $\langle k_i, k_j \rangle$ according to EI-Index, and then compute the scores of every relevant maximal r -radius graph according to Equation 9. Finally, we rank the results and return the $top-k$ r -radius Steiner graphs with the highest scores by refining the corresponding r -radius graphs.

For example, recall the query "DISCOVER, Relational, Database, Papakonstantinou" in Example 5. We first retrieve the relevant maximal r -radius graphs based on EI-Index and compute the corresponding overall scores. We then rank them based on such scores, i.e., $\mathcal{G}_5, \mathcal{G}_4, \mathcal{G}_3$. Finally, we identify the r -radius Steiner graphs on top of the corresponding graph partitions.

However, this method needs to first compute the scores for all relevant maximal r -radius graphs and then rank them. This leads to low efficiency in the presence of large numbers of r -radius Steiner graphs. To alleviate the problem, we introduce an effective technique of progressively identifying the $top-k$ answers. Note that there are many studies of effectively retrieving $top-k$ answers from multiple inverted lists, such as Fagin Algorithm (FA) [7] and Threshold Algorithm (TA) [8]. We can employ such techniques to identify the $top-k$ answers on top of our EI-Index as follows: Given a keyword query $\mathcal{K}=\{k_1, k_2, \dots, k_m\}$, we can retrieve the inverted lists composed of relevant r -radius graphs and corresponding scores for every keyword pair according to our EI-Index. Then, we adopt existing algorithms to progressively identify the $top-k$ answers. The advantage of our approach is obvious – we need not discover the structural relationships by traversing the whole graph. Instead, we first materialize such rich relationships into our EI-Index off-line, and then identify the $top-k$ answers according to EI-Index online.

6. EXPERIMENTAL STUDY

We have designed and performed a comprehensive set of experiments to evaluate the search performance of EASE. We employed the datasets of DBLife[¶], DBLP^{||} and IMDB^{**} to evaluate EASE on unstructured, semi-structured and structured data respectively. There were about 10,000 pages in the DBLife dataset, and the raw file of DBLP was about 400MB. IMDB contained about one million anonymous ratings of approximately 3900 movies made by 6040 users. All datasets were deposited as per submission requirement.

[¶] <http://dblfe.cs.wisc.edu/download/surfaceWeb-2007-07-25.tgz>

^{||} <http://dblp.uni-trier.de/xml>

^{**} <http://www.grouplens.org/node/73> (1,000,000 Data Set(5.73 MB))
http://www.grouplens.org/system/files/million-ml-data.tar_0.gz

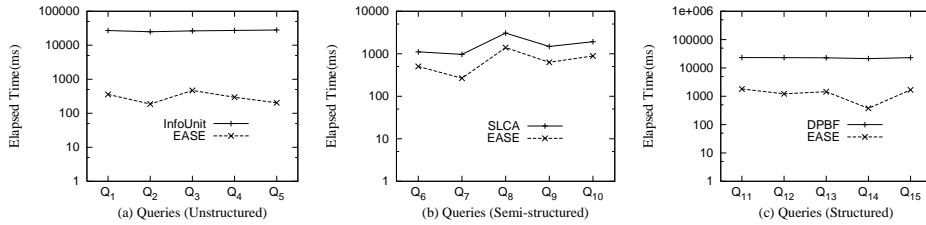


Figure 6: Search efficiency on various queries

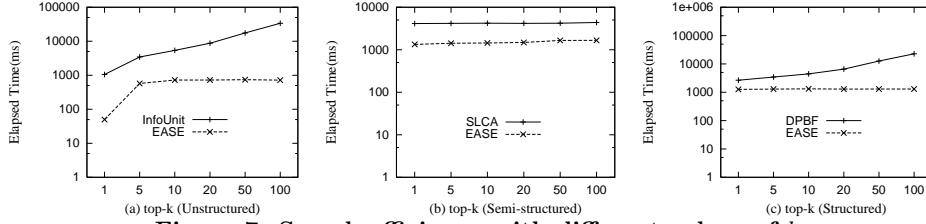


Figure 7: Search efficiency with different values of k

Table 7: Queries employed in the experiments

(a) Queries on unstructured data (DBLife)

Query ID	Queries
Q_1	2007 Conference Data Integration
Q_2	XML Relational Keyword Search
Q_3	Turkey Conference 2007 Uncertainty Fuzziness
Q_4	Berkeley Phone Dataspaces Information Management
Q_5	Beijing Conference 2007 Data Integration

(b) Queries on semi-structured data (DBLP)

Query ID	Queries
Q_6	Information Retrieval Database
Q_7	IR Database
Q_8	DB IR XML
Q_9	XML Relational Keyword Search
Q_{10}	Data Mining Algorithm 2006

(c) Queries on structured data (IMDB)

Query ID	Queries
Q_{11}	Lethal Weapon 4 academic
Q_{12}	Police Academy 3 customer
Q_{13}	Halloween 5 college
Q_{14}	Love 45 tradesman
Q_{15}	Robocop 3 college

(d) Queries on heterogeneous data

Query ID	Queries
Q_{16}	XML Keyword Search 2007
Q_{17}	Dennis Shasha Database Tuning
Q_{18}	Dataspaces 2006 Information Management
Q_{19}	Database Indexing Ranking Search
Q_{20}	Romance Action Educator

The experiments were conducted on an Intel(R) Core(TM) 2.0GHz computer with 2GB of RAM running Windows Vista, and the algorithms were implemented in Java. We used MYSQL 5.0.45^{††} to maintain the graph index. We employed the threshold algorithm [8] to progressively identify the $top-k$ answers based on the *EI-Index*. We compared our approach with existing state-of-the-art approaches. For the unstructured data, we compared EASE with DBLife[5] by submitting keyword queries to its interface and InfoUnit[18]. For the semi-structured data, we compared EASE with SLCA [28] while for structured data, we compared it with DPBF [6]. We selected 20 queries as illustrated in Table 7 for testing different aspects of search problems such as the ability of the methods in capturing data lineage and relationships, and their search accuracy.

The graph w.r.t. DBLife contains about 10,000 nodes, which are far fewer than the approximate 1,000,000 nodes of IMDB and 12,000,000 nodes of DBLP. Note that the graph

w.r.t. IMDB is much denser than that of DBLP, which in turn is denser than that of DBLife. The elapsed time of indexing DBLife, IMDB and DBLP is respectively 13, 25 and 87 minutes. The sizes of the graph indices of DBLife, IMDB and DBLP are respectively 128MB, 96MB and 922MB, compared with the data sizes of 131MB, 30MB and 405MB. The graph index of DBLife is smaller than its data size due to: *i*) tokenization for removing tags, and *ii*) its sparser structure.

6.1 Search Efficiency on Different Datasets

We evaluate the search efficiency of EASE in this section. We tested every algorithm on the selected queries, identified the $top-100$ results and compared their corresponding elapsed time. Figure 6 summarizes the experimental results. We observe that EASE always achieves the best performance for various keyword queries and outperforms the other methods significantly on different datasets.

On unstructured data, InfoUnit identifies Steiner trees by traversing the whole graph to discover structural relationships online, which is not as efficient as our method that materializes rich structural information into the graph indices for facilitating online keyword-based query processing.

On semi-structured data, SLCA has to first identify all the relevant results, and then ranks them and returns the $top-k$ answers with the highest scores. In contrast, EASE employs the threshold based algorithm [8] to progressively identify $top-k$ answers and thus produces a dramatic improvement over the existing methods.

On structured data, although DPBF can progressively compute the $top-k$ answers, it has to discover the structural information between tuples in different relational tables online. At the same time, it also needs to identify the Steiner trees on top of the whole graph. These needs lead to low efficiency. In contrast, EASE first identifies the r -radius graphs based on *EI-Index* and then constructs r -radius Steiner graphs on top of the corresponding graph partitions. These partitions are much smaller than the whole graph, and the method therefore achieves better search performance and outperforms DPBF significantly.

To better evaluate the selected algorithms, we identified the $top-k$ answers with different values of k and compared the corresponding average elapsed time. Figure 7 illustrates the experimental results obtained. We observe that EASE outperforms the other methods significantly. We note that, with the increase of k values, the elapsed time of InfoUnit

^{††}<http://dev.mysql.com/downloads/mysql/5.0.html>

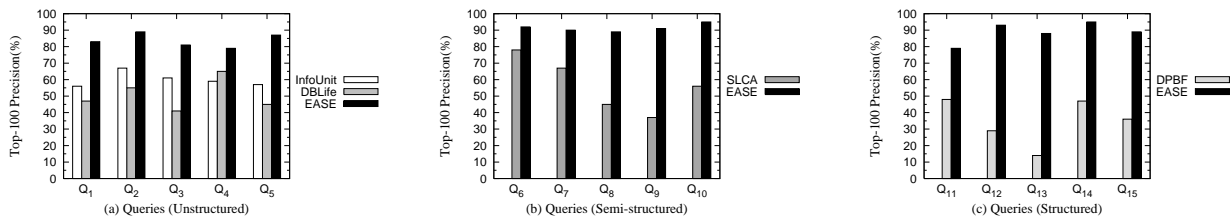


Figure 8: top -100 precision on various queries

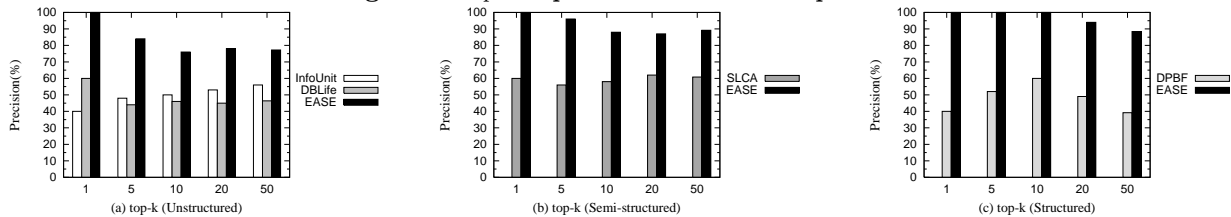


Figure 9: top - k precision with different values of k

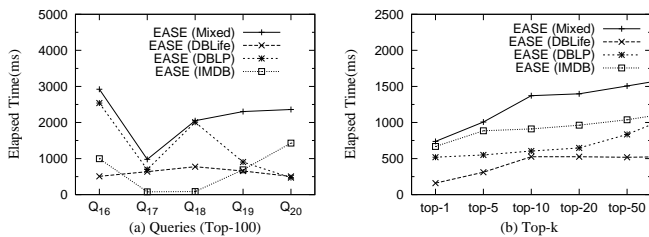


Figure 10: Search efficiency on heterogenous data

and DPBF also increases, but the elapsed time of EASE varies only slightly (as SLCA cannot progressively identify answers, its elapsed time also varies little). The elapsed time of InfoUnit and DPBF to return top -100 results is always a little more than that of top -5 while the elapsed time of EASE remains about the same. This is so because we adopt the threshold based technique to progressively identify the top - k answers.

6.2 Search Efficiency on Heterogenous Datasets

To evaluate the overall performance of EASE on heterogenous data, we mixed the three datasets and tested EASE on the combined data. We constructed the graph index on top of the heterogenous data and employed Q_{16} - Q_{20} to evaluate the search efficiency of EASE over heterogenous data. To better understand the performance of EASE, we first evaluated the elapsed time of identifying the top -100 answers for every query. We then varied the values of k to evaluate the average elapsed time. We also provide the elapsed time of EASE on different datasets as a baseline to show the search efficiency of EASE over heterogenous data. The results are summarized in Figure 10. We can see that EASE is still capable of achieving high search efficiency over heterogenous data and is not much worse off than on homogeneous datasets. That is, even for a large volume of heterogenous data, EASE can still effectively identify the top - k answers.

6.3 Search Accuracy on Different Datasets

This section evaluates search accuracy, indicating the fraction of relevant results in the approximate answer that are correct. We employed the metric, top - k precision, which measures the ratio of the number of relevant answers among the first k answers with the highest scores of an algorithm to k . Answer relevance is judged from discussions of researchers in our database group. Interested readers can test our algorithm through our deposited demo as per submis-

sion requirement. Moreover, to evaluate the search accuracy of the selected algorithms, we identified the top -100 results for every query and compared the corresponding top -100 precision. Figure 8 illustrates the experimental results obtained. We observe that EASE achieves much higher precision than existing methods such as, DBLife, InfoUnit, SLCA and DPBF on the corresponding datasets.

On unstructured data, as InfoUnit and EASE integrate relevant pages to answer keyword queries, they achieve much higher precision than DBLife. EASE is better than InfoUnit because EASE employs a ranking mechanism that takes into account both the structural compactness of answers from the DB viewpoint and textual relevancy from the IR viewpoint. On semi-structured data, r -radius Steiner graphs, alongside content nodes, also contain some relevant elements, such as the Steiner nodes, which may expand the answerability. They are more meaningful than the subtrees/path-trees rooted at LCAs (or its variants) of existing methods, which may miss some relevant information. Thus, EASE yields higher precision. On structured data, EASE also outperforms DPBF in that: (i) r -radius Steiner graphs are more meaningful than Steiner trees; and (ii) Again, this is due to the consideration of structural compactness between input keywords in our ranking function.

For example, consider query Q_7 . Its answer should be the papers about “Database” and “IR”. However, “IR” may appear in the name of an author, and traditional methods such as SLCA cannot distinguish the papers entitled with “Database” and “IR” and the papers with “IR” being a term of the author and “Database” being a term of the title. They may mistakenly take the latter as answers, leading to low search accuracy. As another example, two input keywords may appear in two different authors of a paper, and the conventional methods will mistakenly take the two keywords as an author. EASE ranks the results with structures that are less compact lower. This is done by means of our structural relevancy based ranking method. Hence, EASE achieves much higher precision.

As users are usually interested in top - k answers, we varied different values of k to evaluate the selected algorithms. The average results of the top - k precision for Q_1 - Q_{15} are illustrated in Figure 9. As expected, EASE consistently achieved high precision in many queries, which is approximately 10-30% higher than those of InfoUnit, DBLife, DPBF and SLCA.

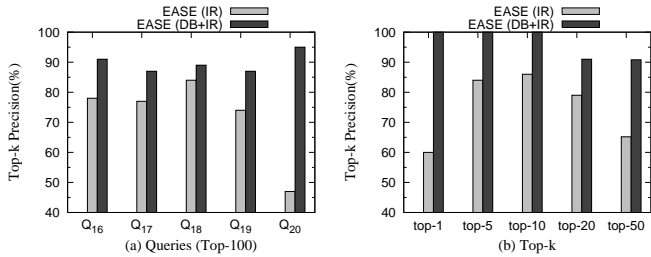


Figure 11: Search accuracy on heterogeneous data

6.4 Search Accuracy on Heterogeneous Datasets

To evaluate the robustness of our algorithm, we evaluate its search accuracy using the heterogeneous dataset. We first evaluated search accuracy by identifying the *top-100* results on various queries, and then varied different values of k to evaluate the average precision. For a better understanding of our ranking method, we compared EASE employing IR ranking parameters (ref. Equation 5) with EASE considering both DB and IR (ref. Equation 9). Figure 11 summarizes the experimental results. We observe that EASE(DB+IR) consistently achieves much higher accuracy than EASE(IR). This again confirms that TF-IDF based IR ranking cannot capture the structural relationships effectively.

For example, consider query Q_{20} . Some results on DBLife dataset will obtain higher IR scores than those on IMDB dataset as the term frequencies of some documents on DBLife dataset are larger than those on IMDB dataset. Thus, the IR based ranking method will rank the results on DBLife dataset higher. However, in the IMDB dataset, which is a collection of ratings of movies made by users, the three input keywords can describe the instances of users with occupation of Educator scoring the movies with genres of Romance and Action. Hence, the three keywords are more compact and highly relevant to IMDB dataset. Our DB+IR ranking method considers the structural relevancy and ranks the results on the IMDB dataset higher, and thus achieves much higher precision. This comparison reflects the effectiveness and applicability of our proposed ranking mechanism.

6.5 Sensitivity of Indexing

In this section, we evaluate the sensitivity of our indexing method. We first created the indices on top of the DBLife dataset, and then inserted 5% additional new irrelevant data in each step. After each set of insertions, we evaluated EASE on the updated indices by using queries Q_1 - Q_5 . Figure 12 summarizes the effects of insertions on our indexing method. We observe that the search accuracy of EASE does not degrade drastically with the increase of new data. This demonstrates the robustness and scalability of our method with respect to data insertions.

7. RELATED WORK

The first area of research related to our work is keyword search over relational databases by identifying Steiner trees. As opposed to the traditional Steiner tree based methods, which identify the structural relationships online, EASE identifies and materializes the rather rich structural relationships so as to improve the online processing of keyword queries.

DBXplorer[1], DISCOVER-I[15], DISCOVER-II[13], BANKS-I [2] and BANKS-II [17] are systems built on top of relational databases. DISCOVER and DBXplorer generate trees of tuples connected through primary-foreign key relationships that contain all of the input keywords. BANKS

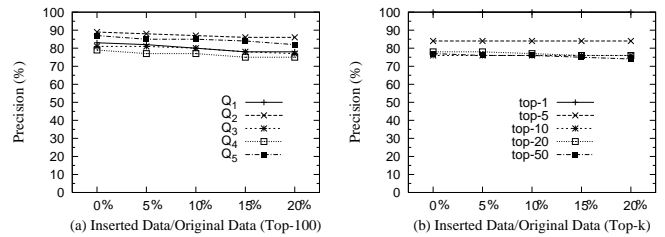


Figure 12: Sensitivity of indexing

identifies connected trees in a labeled graph by using an approximation of the Steiner tree problem. DISCOVER-II considers the problem of keyword proximity search in terms of disjunctive semantics, as opposed to DISCOVER-I which only considers conjunctive semantics. Kacholia et al. [17] presented the bidirectional strategy (BANKS-II) to improve the efficiency of keyword search over graph data. However, their method still works by identifying Steiner trees from the whole graph, which is inefficient as it is rather difficult to identify structural relationships through inverted indices. Liu et al. [20] proposed a novel ranking strategy to solve the effectiveness problem for relational databases. It employs phrase-based and concept-based models to improve search effectiveness by introducing IR techniques.

More recently, Ding et al. [6] employed dynamic programming (DPBF) to improve the efficiency of identifying Steiner trees. Guo et al. [10] proposed data topology search to retrieve meaningful structures from richer structural data, such as complex biological databases. He et al. [12] proposed a partition-based method to improve search efficiency with a novel BLINKS index. Markowitz et al. [23] studied the problem of keyword search over relational data streams in a first attempt to answer keyword search over relational data streams. Luo et al. [22] proposed a new ranking method that adapts state-of-the-art IR ranking function and principles into the ranking trees of joined database tuples. In addition, Yu et al. [29] studied the problem of relational data source selection in P2P environments by summarizing the relationships between keywords in underlying databases.

In terms of keyword search over XML documents, the subtrees rooted at the lowest common ancestors (LCAs) of content nodes have been proposed as answers. As an extension of LCA, SLCA, Multiway-SLCA, XSeek and GDMCT have recently been proposed to answer keyword queries over XML documents in [28], [25], [21] and [14] respectively. SLCA [28] can avoid the false positives of LCA but it does so at the expense of false negatives. Multiway-SLCA [25] offers a search paradigm in support of keyword search beyond the traditional AND semantics, including both AND and OR boolean operators. GDMCT [14] returns grouped connected trees as answers but it still needs to traverse the whole graph to identify answers by employing the inverted index. XSeek [21] generates return nodes which can be explicitly inferred from keywords or dynamically constructed according to entities in the data that are relevant to the search.

XRANK [11] and XSearch [4] are systems facilitating keyword search for XML documents, and they return connected subtrees as answers for keyword queries. XRANK presents a ranking method, where for a given tree \mathcal{T} containing all the keywords, a score is assigned to \mathcal{T} with an adaptation of PageRank for XML documents. XSearch focuses on semantics and the ranking of results; during execution, it uses an all-pairs interconnection index to check the connectivity

between nodes. XKeyword [16] is a system that offers keyword proximity search over XML documents that conform to an XML schema. However, it needs to compute candidate networks and thus is constrained by schemas. TopX [26] is a prototype search engine for the ranked retrieval of XML, but it processes XML queries with support for XPath axes but not the more simple keyword queries, and it cannot adapt to relational databases. Graupmann et al. [9] presented the SphereSearch engine to provide unified ranked retrieval on heterogeneous XML and Web data. However, it is orthogonal to our method in that: *i*) it does not support relational databases and it transforms HTML documents into XML, and *ii*) it depends on its own query language to discover structural relationships and thus is not a pure keyword based search method. Chaudhuri, Ramakrishnan and Weikum [3] point out a number of interesting research opportunities for integrating DB and IR technologies. In [27], Weikum provides a summary of the existing DB and IR techniques and discusses the difficulties, opportunities and challenges of combining DB and IR.

8. CONCLUSION

In this paper, we have proposed an efficient and adaptive keyword search method, EASE, to answer keyword queries over unstructured, semi-structured and structured data. EASE seamlessly integrates the efficient query evaluation of DB and the adaptive scoring models of IR for the ranking of results. EASE models heterogeneous data as graphs and processes keyword queries on the graphs. To the best of our knowledge, this is the first attempt to efficiently and adaptively process keyword queries on heterogeneous data. We have proposed summarizing and clustering the graphs, and devised effective graph indices to materialize structural relationships for fast and accurate response. To facilitate efficient keyword-based query processing, we have examined the issues of indexing and ranking by taking into account both the structural compactness of r -radius graphs from the DB point of view and textual relevancy from the IR viewpoint. Finally, we have conducted an extensive performance study to evaluate the efficiency and effectiveness of our approach using real datasets. The experimental results show that EASE achieves both high search efficiency and quality for keyword search over heterogeneous data, and significantly outperforms the existing methods.

9. ACKNOWLEDGEMENT

The research of B.C. Ooi was in part funded by NUS Grant R-252-000-338-112. This work is in part supported by the National Natural Science Foundation of China under Grant No.60573094, the National High Technology Development 863 Program of China under Grant No.2007AA01Z152 and 2006AA01A101, the National Grand Fundamental Research 973 Program of China under Grant No.2006CB303103.

10. REPEATABILITY ASSESSMENT RESULT

All the results in this paper were verified by the SIGMOD repeatability committee. Code and data used in the paper are available at <http://www.sigmod.org/codearchive/sigmod2008/>.

11. REFERENCES

- [1] S. Agrawal, S. Chaudhuri, and G. Das. Dbxplorer: A system for keyword-based search over relational databases. In *ICDE*, pages 5–16, 2002.
- [2] G. Bhalotia, A. Hulgeri, C. Nakhe, S. Chakrabarti, and S. Sudarshan. Keyword searching and browsing in databases using banks. In *ICDE*, pages 431–440, 2002.
- [3] S. Chaudhuri, R. Ramakrishnan, and G. Weikum. Integrating db and ir technologies: What is the sound of one hand clapping? In *CIDR*, pages 1–12, 2005.
- [4] S. Cohen, J. Mamou, Y. Kanza, and Y. Sagiv. Xsearch: A semantic search engine for XML. In *VLDB*, 2003.
- [5] P. DeRose, W. Shen, F. Chen, Y. Lee, D. Burdick, A. Doan, and R. Ramakrishnan. DbLife: A community information management platform for the database research community. In *CIDR*, 2007.
- [6] B. Ding, J. X. Yu, S. Wang, L. Qin, X. Zhang, and X. Lin. Finding top-k min-cost connected trees in databases. In *ICDE*, 2007.
- [7] R. Fagin. Combining fuzzy information from multiple systems. In *PODS*, pages 216–226, 1996.
- [8] R. Fagin. Fuzzy queries in multimedia database systems. In *PODS*, pages 1–10, 1998.
- [9] J. Graupmann, R. Schenkel, and G. Weikum. The spheresearch engine for unified ranked retrieval of heterogeneous XML and web documents. In *VLDB*, pages 529–540, 2005.
- [10] L. Guo, J. Shanmugasundaram, and G. Yona. Topology search over biological databases. In *ICDE*, 2007.
- [11] L. Guo, F. Shao, C. Botev, and J. Shanmugasundaram. Xrank: Ranked keyword search over XML documents. In *SIGMOD*, pages 16–27, 2003.
- [12] H. He, H. Wang, J. Yang, and P. Yu. Blinks: Ranked keyword searches on graphs. In *SIGMOD*, 2007.
- [13] V. Hristidis, L. Gravano, and Y. Papakonstantinou. Efficient ir-style keyword search over relational databases. In *VLDB*, pages 850–861, 2003.
- [14] V. Hristidis, N. Koudas, Y. Papakonstantinou, and D. Srivastava. Keyword proximity search in XML trees. In *IEEE TKDE 18(4)*, pages 525–539, 2006.
- [15] V. Hristidis and Y. Papakonstantinou. Discover: Keyword search in relational databases. In *VLDB*, 2002.
- [16] V. Hristidis, Y. Papakonstantinou, and A. Balmin. Keyword proximity search on XML graphs. In *ICDE*, pages 367–378, 2003.
- [17] V. Kacholia, S. Pandit, S. Chakrabarti, S. Sudarshan, R. Desai, and H. Karambelkar. Bidirectional expansion for keyword search on graph databases. In *VLDB*, pages 505–516, 2005.
- [18] W.-S. Li, K. S. Candan, Q. Vu, and D. Agrawal. Retrieving and organizing web pages by ‘information unit’. In *WWW*, 2001.
- [19] G. Li, J. Feng, J. Wang, and L. Zhou. Efficient Keyword Search for Valuable LCAs over XML Documents. In *CIKM*, 2007.
- [20] F. Liu, C. Yu, W. Meng, and A. Chowdhury. Effective keyword search in relational databases. In *SIGMOD*, 2006.
- [21] Z. Liu and Y. Chen. Identifying return information for XML keyword search. In *SIGMOD*, 2007.
- [22] Y. Luo, X. Lin, W. Wang, and X. Zhou. Spark: Top-k keyword query in relational databases. In *SIGMOD*, 2007.
- [23] A. Markowetz, Y. Yang, and D. Papadias. Keyword search on relational data streams. In *SIGMOD*, 2007.
- [24] M. Mutsuzaki, M. Theobald, A. Keijzer, J. Widom, P. Agrawal, and et al. Trio-one: Layering uncertainty and lineage on a conventional dbms. In *CIDR*, 2007.
- [25] C. Sun, C. Y. Chan, and A. K. Goenka. Multiway sica-based keyword search in XML data. In *WWW*, 2007.
- [26] M. Theobald, R. Schenkel, and G. Weikum. An efficient and versatile query engine for topx search. In *VLDB*, 2005.
- [27] G. Weikum. Db & ir: Both sides now (keynote). In *SIGMOD*, 2007.
- [28] Y. Xu and Y. Papakonstantinou. Efficient keyword search for smallest lcas in XML databases. In *SIGMOD*, pages 527–538, 2005.
- [29] B. Yu, G. Li, K. Sollins, and A. K. H. Tung. Effective keyword-based selection of relational databases. In *SIGMOD*, 2007.