

Parse Tree Database for Information Extraction

Luis Tari, Phan Huy Tu, Jörg Hakenberg, Yi Chen, Tran Cao Son, Graciela Gonzalez and Chitta Baral

Abstract—Information extraction systems are traditionally implemented as a pipeline of special-purpose processing modules targeting the extraction of a particular kind of information. A major drawback of such approach is that whenever a new extraction goal emerges or a module is improved, extraction has to be re-applied from scratch to the entire text corpus even though only a small part of the corpus might be affected. In this paper, we describe a novel approach for information extraction so that extraction needs are expressed in the form of database queries, which are evaluated and optimized by databases. Using database queries for information extraction enables generic extraction and minimizes reprocessing of data. In addition, our approach provides two different query generation components that can automatically form database queries for extraction from training datasets, as well as from unlabeled data through a mechanism inspired by the pseudo-relevance feedback approach found in protein-protein interactions and drug-protein-metabolic relations from two sets of corpus. Experiments show that our approach achieves a precision of 83.6% and recall of 58.6% (F-measure of 64.2%) for the extraction of protein-protein interactions from the BioCreative 2 corpus, while achieving a precision of 85.0% and recall of 26.0% (F-measure of 39.8%) for drug-protein-metabolic relations.

Index Terms—Text mining, Query languages, Information Storage and Retrieval

1 INTRODUCTION

IT is estimated that each year more than 600,000 articles are published in the biomedical literature, with about 19 million publication entries being stored in the Medline database¹. To uncover information from such a large corpus of documents, it is vital to address the need in an automated manner. The field of information extraction (IE) seeks to develop methods for fetching concise, structured information from natural language text. Examples of such structured information are the extraction of entities and relationships between entities.

IE is typically seen as a one-time process for the extraction of a particular kind of relationships of interest from a document collection. IE is usually deployed as a pipeline of special-purpose programs, which include sentence splitters, tokenizers, named entity recognizers, shallow or deep syntactic parsers, and extraction based on a collection of patterns. The large demand of IE in various domain results in the development of frameworks such as UIMA [1] and GATE [2], providing a way to perform extraction by defining workflows of components. Such kind of extraction frameworks are usually file-based so that large amount of processed data can be utilized between components. In this setting, relational databases play a limited role of storing the extracted relationships.

While such framework is suitable for one-time extraction, it is important to notice that there are cases when IE has to be performed repeatedly even on the same document collection. Examples of such cases include the processing of web documents with modified content [3], the availability of updated ontologies or improved components for named entity recognition, and the realization of new target relationships for extraction. Using the existing extraction frameworks in any of these scenarios means that it is necessary to reprocess the entire text collection, which can be computationally expensive. It is important to realize that it is unnecessary to reprocess the entire collection when certain components in the pipeline or extraction goals are changed. Consider a biology-oriented scenario when the original goal of extracting protein-protein interactions changes into finding gene-disease associations from the same corpus of text. Consider another scenario where the extraction goal remains the same, but an updated ontology or an improved model based on statistical learning approach becomes available for named entity recognition. Changes in these scenarios typically only affect a portion of the text corpus. Thus an ideal framework needs to have the capability of managing processed data and performing *incremental extraction* to identify which part of the data is affected by the change of components or goals.

In this paper, we propose a new paradigm for information extraction that utilizes database management systems as an essential component of our extraction framework. Database management systems become a logical framework of choice that can serve such dynamic extraction needs over file-based storage systems. As illustrated in Figure 1, text processing components such as named entity recognizers and syntactic parsers are deployed for the entire collection. The intermediate output of the processing modules is stored in a relational database known as the *parse tree database*. Extraction then

- L. Tari, P.H. Tu, J. Hakenberg, Y. Chen and C. Baral are with the Department of Computer Science and Engineering, Arizona State University, Tempe, AZ, 85287.
E-mail: luis.tari@asu.edu
- G. Gonzalez is with the Department of Biomedical Informatics, Arizona State University, Phoenix, AZ 85004.
- T. C. Son is with the Department of Computer Science, New Mexico State Univ., Las Cruces, NM 88003.

1. Medline statistics:
http://www.nlm.nih.gov/bsd/licensee/2009_stats/2009_Totals.html

becomes a matter of issuing database queries in the form of *parse tree query language* (PTQL). In the event of a change of extraction goals or a module update, the responsible module is deployed for the entire text corpus and the processed data is populated into the parse tree database with the previously processed data. Incremental extraction is performed so that database queries are issued to identify sentences with newly recognized mentions. Once the affected sentences are identified, extraction can then be performed only on such sentences rather than the entire corpus. By storing the processed data, our approach avoids the need to reprocess the entire collection of text unlike the file-based pipeline approaches. Avoiding reprocessing of data is particularly important for extraction in the biomedical domain, where a full processing of all 17 million Medline abstracts took about more than 36K hours of CPU time using a single-core CPU with 2-GHz and 2 GB of RAM. In this case, the Link Grammar parser [4] contributes to a large portion of the time spent in text processing.




Fig. 1. An overview of our extraction framework illustrating how the parse tree database is used for extraction and update of modules.

We highlight the technical contributions of the architecture proposed in this paper.

- *Novel Database-Centric Framework for Information Extraction.* Unlike traditional approaches for IE, our approach is to store intermediate text processing output in a specialized database called the *parse tree database*. Extraction is formulated as queries so that it is no longer necessary to write and run special-purpose programs for each specific extraction goal. Our approach minimizes the need of reprocessing the entire collection of text in the presence of new extraction goals and deployment of improved processing components.
- *Query Language for Information Extraction.* The use of grammatical structures such as constituent trees and linkages to define extraction patterns requires traversals of paths in constituent trees, as well as links and link types between node pairs during evaluation of extraction patterns. As query languages such as XPath and XQuery are not suitable for extracting linguistic patterns [5], we designed and implemented a query language called *parse tree query language* (PTQL) for pattern extraction which effectively achieves diverse IE goals.
- *Automated Query Generation.* Manually writing extraction queries can be a time-consuming, labor-intensive process, and such approach is likely to achieve unsat-

isfactory extraction performance. Our system provides two forms of automated query generation either in the presence or absence of training data.

Our information extraction framework is composed of two phases:

- *Initial Phase:* we perform a one-time parse, entity recognition and tagging (identifying individual entries as belonging to a class of interest) on the whole corpus based on current knowledge. The generated syntactic parse trees and semantic entity tagging of the processed text is stored in a *parse tree database* (PTDB).
- *Extraction Phase:* Extracting particular kinds of relations can be done by issuing an appropriate query to the PTDB. To express extraction patterns, we designed and implemented a query language called *parse tree query language* (PTQL) that is suitable for generic extraction. Our system not only allows a user to issue PTQL queries for extraction, but it can also automatically generate queries from training data or user keyword-based queries.

The rest of the paper is organized as follows. We first present the necessary background in Section 2.2. In Section 3, the system architecture of our extraction framework is discussed in details, which includes the PTDB, the query language PTQL and its evaluation. We then describe the two query generation components in our framework in Section 4, which enable the generation of extraction queries from both labeled and unlabeled data. The query performance of our approach and the quality of the extracted results are presented in Section 5. We describe the related work and conclude in Sections 6 and 7.

2 BACKGROUND

2.1 Information extraction

IE has been an active research area that seeks techniques to uncover information from a large collection of text. Examples of common IE tasks include the identification of entities (such as protein names, disease names), extraction of relationships between entities (such as interactions between a pair of proteins) and extraction of entity attributes (such as coreference resolution that identifies different variants of mentions corresponding to the same entity) from text. Readers are referred to [6] for a recent and comprehensive survey of IE.

The examples and experiments used in our paper involve the use of grammatical structures for relationship extraction. Cooccurrences of entities is a typical method in relationship extraction, but often leads to imprecise results. Consider that our goal is to extract relations between drug and proteins from the following sentence:

Quetiapine is metabolised by CYP3A4 and sertindole by CYP2D6. (PMID:10422890)

By utilizing our grammatical knowledge, a human reader can observe that $\langle \text{CYP3A4, metabolise, quetiapine} \rangle$ and $\langle \text{CYP2D6, metabolise, sertindole} \rangle$ are the only correct triplet relations for the above sentence. However, if we consider cooccurrences of entities as a

criteria to extract relationships, incorrect relationships such as $\langle \text{CYP3A4, metabolise, sertindole} \rangle$ and $\langle \text{CYP2D6, metabolise, quetiapine} \rangle$ would also be extracted from the above sentence. This simple example highlights the need of grammatical knowledge in performing relationship extraction.

A typical IE setting involves a pipeline of text processing modules in order to perform relationship extraction. These include:

- *sentence splitting*: identifies sentences from a paragraph of text
- *tokenization*: identifies word tokens from sentences
- *named entity recognition*: identifies mentions of entity types of interest
- *syntactic parsing*: identifies grammatical structures of sentences
- *pattern matching*: obtains relationships based on a set of extraction patterns that utilize lexical, syntactic and semantic features

Figure 2 illustrates a typical text processing workflow in order to perform extraction of relationships. Extraction patterns are typically obtained through manually written patterns compiled by experts or automatically generated patterns based on training data. Different kinds of parsers, which include shallow and deep parsers, can be utilized in the pipeline. In our work, the Link Grammar parser [7] is utilized as part of our extraction approach. We describe the basic terminologies involved in Link Grammar in the next subsection.

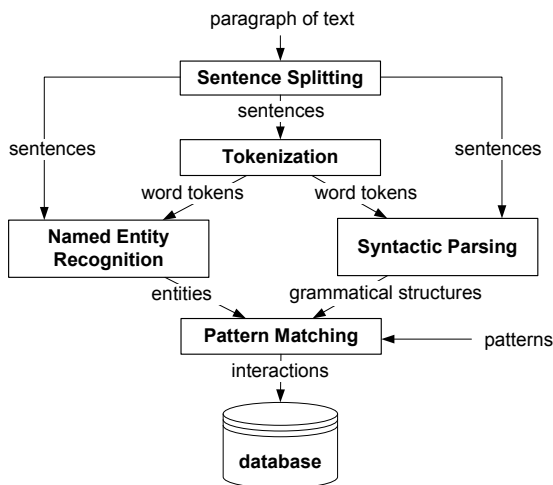


Fig. 2. A workflow of text processing modules that takes a paragraph of text as input to perform relationship extraction.

2.2 Link Grammar

The Link Grammar parser is a dependency parser based on the Link Grammar theory [7]. Link Grammar consists of a set of words and linking requirements between words. A *sentence* of the language is defined as a sequence of words such that the links connecting

the words satisfy the following properties: (i) the links do not cross, (ii) the words form a connected graph, and (iii) the links satisfy the linking requirements of each word in the sentence. The output of the parser, called a *linkage*, shows the dependencies between pairs of words in the sentence. Figure 3 shows an example for the sentence “RAD53, which activates DNA damage, positively regulates the DBF4 protein” (PMID:10049915). The linkage contains several links, which include link S connecting the subject-noun RAD53 to the transitive verb regulates, the O link connecting the transitive verb regulates to the direct object DBF4 and the MX*r link connecting the relative pronoun which to the noun RAD53. For a complete description of links, we refer the reader to [4].




Fig. 3. Linkage of the sentence “RAD53, which activates DNA damage, positively regulates the DBF4 protein”, showing only the main link types in the linkage.

Besides producing linkages, the Link Grammar parser is also capable of outputting constituent trees. A *constituent tree* is a syntactic tree of a sentence with the nodes represented by part-of-speech tags and words of the sentences in the leaf nodes. For instance, the corresponding constituent tree for the above sentence is illustrated in Figure 4. In the constituent tree, S stands for a sentence/clause, SBAR for a clause containing a clause with a relative pronoun, WHNP for a clause containing a relative pronoun, NP for a noun phrase, VP for a verb phrase and ADVP for an adverb phrase. The leaf nodes of the constituent tree represent the words of the sentence and their part-of-speech tags. For words that are not recognizable by the parser, the tag U is given for such words for unknown part-of-speeches.

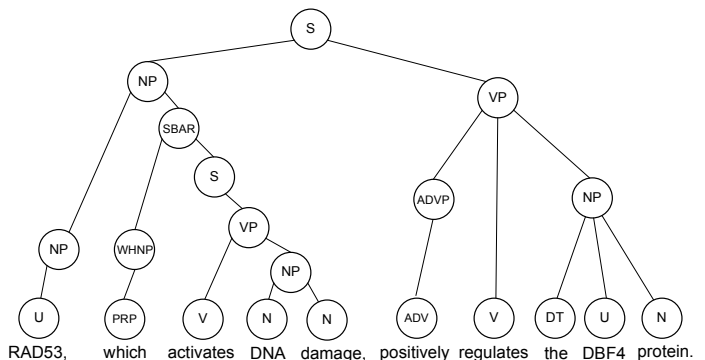


Fig. 4. Constituent tree of the sentence “RAD53, which activates DNA damage, positively regulates the DBF4 protein”

3 SYSTEM ARCHITECTURE

We first give an overview of our approach, and discuss each of the major components of our system in this section. Our approach is composed of two phases: *initial phase* for processing of text and *extraction phase* for using database queries to perform extraction. The *Text Processor* in the initial phase is responsible for corpus processing and stores the processed information in the *Parse Tree Database (PTDB)*. The extraction patterns over parse trees can be expressed in our proposed *parse tree query language (PTQL)*. In the extraction phase, the *PTQL query evaluator* takes a PTQL query and transforms it into keyword-based queries and SQL queries, which are evaluated by the underlying RDBMS and IR engine. To speed up query evaluation, the *index builder* creates an inverted index for the indexing of sentences according to words and the corresponding entity types. Figure 5 illustrates the system architecture of our approach.

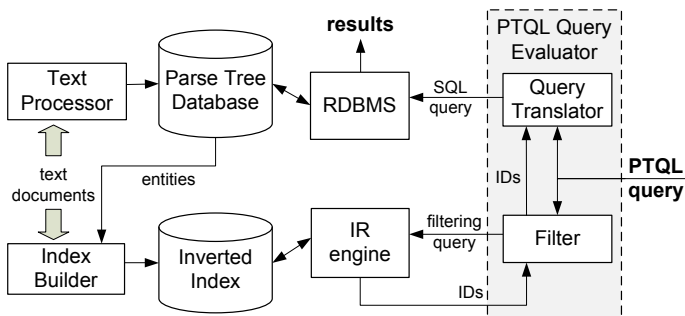


Fig. 5. System architecture of the PTQL framework. The framework includes the parse tree database for storing intermediate processed information and the query evaluator for the evaluation of PTQL queries through filtering and translation to SQL queries.

Our approach provides two modes of generating PTQL queries for the purpose of information extraction: *training set driven query generation* and *pseudo-relevance feedback driven query generation*. To generate a set of patterns for information extraction using the training set driven approach, the *pattern generator* first automatically annotates an unlabeled document collection with information drawn from a problem-specific database. This step necessitates a method for precise recognition and normalization of protein mentions. From this labeled data, initial phrases referring to interactions are extracted. These phrases are then refined to compute consensus patterns and the resulting PTQL queries are generated by the *query generator*. However, training data is not always readily available for certain relationships due to the inherent cost of creating a training corpus. In that regards, our approach provides the pseudo-relevance feedback driven approach that takes keyword-based queries, and the *PTQL query generator* then finds common grammatical patterns among the top- k retrieved sentences to generate PTQL queries.

We first describe the parse tree database and the

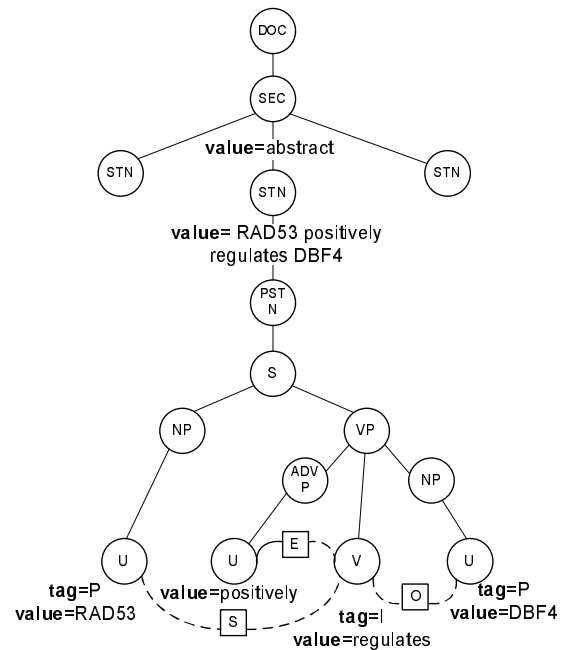


Fig. 6. An example of a parse tree for a document, which includes sections of the document, sentences and the corresponding parse trees. The attribute *Tag* indicates the semantic type of a word, in which P stands for protein names and I for interaction words.

syntax of PTQL before we provide details of how PTQL queries are processed.

3.1 Parse tree database and inverted index

The *Text Processor* parses Medline abstracts with the Link Grammar parser [4], and identifies entities in the sentences using BANNER [8] to recognize gene/protein names and MetaMap [9] to recognize other entity types that include disease and drug names. Each document is represented as a hierarchical representation called the *parse tree of a document*, and the parse trees of all documents in the document collection constitute the *parse tree database (PTDB)*. The detailed schema and its description can be found in Appendix A. A *parse tree* is composed of a constituent tree and a linkage. A *constituent tree* is a syntactic tree of a sentence with the nodes represented by part-of-speech tags and leafs corresponding to words in the sentence. A *linkage*, on the other hand, represents the syntactic dependencies (or links) between pairs of words in a sentence. Each node in the parse tree has labels and attributes capturing the document structure (such as title, sections, sentences), part-of-speech tags, and entity types of corresponding words.

Figure 6 shows an example of a parse tree for a Medline abstract. The parse tree contains the root node labeled as *DOC* and each node represents an element in the document which can be a section (*SEC*), a sentence (*STN*), or a parse tree for a sentence (*PSTN*). A node labeled as *STN* may have more than one child labeled with *PSTN* to allow the storage of multiple parse trees.

The node below the PSTN node indicates the start of the *parse tree*, which includes the constituent tree and linkage of the sentence. A solid line represents a parent-child relationship between two nodes in the constituent tree, whereas a dotted line represents a link between two words of the sentence. In the constituent tree, nodes S, NP, VP and ADVP stand for a sentence, a noun phrase, a verb phrase and an adverb phrase respectively. The linkage contains three different links: the S link connects the subject-noun RAD53 to the transitive verb regulates, the O link connects the transitive verb regulates to the direct object DBF4 and the E link connects the verb-modifying adverb positively to the verb regulates. The square box on a dotted line indicates the link type between two words. Each leaf node in a parse tree has value and tag attributes. The value attribute stores the text representation of a node, while the tag attribute indicates the entity type of a leaf node. For instance, a protein is marked with a tag P, a drug name with a tag D and an interaction word is marked with I.

Another essential component of our system architecture is an *inverted index* maintained by an IR engine such as Lucene². This inverted index enables the efficient processing of PTQL queries, which will be discussed in details in a later section. As illustrated in Figure 7, the *index builder* relies on a text preprocessor to recognize entities and replace the entities with identifiers in the sentences. Each sentence in the documents are indexed on its own so that each keyword-based filtering query retrieves a sentence rather than the entire document. Assuming that the *concepts* of interest are entities such as protein and drug names, and interaction verbs in the form of present tense and past participle. We use the identifiers DRUGNAME, PROTNAME, IVERB-S and IVERB-D to represent these concepts respectively. The index builder includes the original sentences in the inverted index, as well as sentences with entities replaced with identifiers. For instance, the sentence “RAD53 positively regulates DBF4” is indexed as PROTNAME positively IVERB-S PROTNAME under the field name sent-proc in the inverted index. The approach of indexing sentences with replaced identifiers is similar to [10], [11], [12]. Unlike [10], our approach requires no modification to the structure of the inverted index in order to process variabilized queries. For efficient access of the hierarchical structure and horizontal relations among nodes, we adopt the labeling scheme used in LPath [13], [5] and present the scheme in Appendix B.

3.2 PTQL: Parse Tree Query Language

A fundamental design criteria for the query language is the ability of expressing linguistic patterns based on constituent trees. Standard XML query languages such as XPath [14] and XQuery [15] seem to be the ideal candidates for querying parse trees. However, the inability of

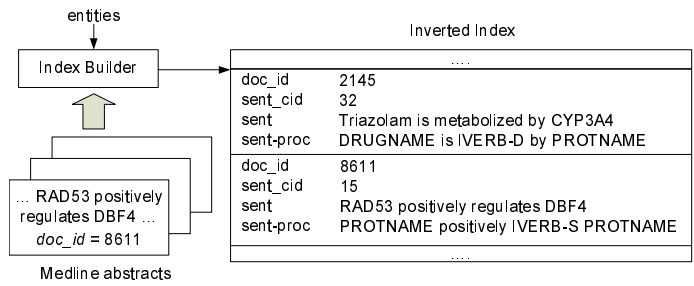


Fig. 7. An extended inverted index to handle queries that involve concepts rather than just instances

expressing immediate-following siblings and immediate-preceding siblings in these standard XML query languages, as shown in [16], leads to the development of LPath [13], [5] as a query language for linguistic queries on constituent trees. An additional design criteria for the query language is the ability to express linguistic patterns based on dependency grammar, such as Link Grammar [7]. Links and link types can be useful in linguistic patterns, such as the type MXsr connects a relative pronoun to its corresponding noun. However, languages such as XQuery and LPath can only express ancestor-descendant and sibling relations between nodes. One of the novel features of our proposed query language PTQL is the ability to express links and link types between pairs of nodes, so that PTQL can be used to express linguistic patterns based on constituent trees and links, as well as link types.

We propose a high level extraction query language called PTQL. PTQL is an extension of the linguistic query language LPath [13], [5] that allows queries to be performed not only on the constituent trees but also the syntactic links between words on linkages. A PTQL query is made up of four components: (i) tree patterns, (ii) link conditions, (iii) proximity conditions, and (iv) return expression. A *tree pattern* describes the hierarchical structure and the horizontal order between the nodes of the parse tree. A *link condition* describes the linking requirements between nodes, while a *proximity condition* is to find words that are within a specified number of words. A *return expression* defines what to return. The EBNF grammar for PTQL is shown in Figure 8. Before going into details of the definition of PTQL queries and its usage, we start with the basic element of PTQL queries called node expressions.

Definition 1: A *node expression* is an expression of the form $X[\langle \text{pred exp} \rangle](x)$ where X is a node name, e.g., a sentence (STN), a parse sentence (PSTN), a noun phrase (NP), or ? (a node of any name), $\langle \text{pred exp} \rangle$ (predicate expression) is a boolean formula of expressions of the form $\langle \text{attribute} \rangle \langle \text{op} \rangle \langle \text{value} \rangle$, and x is a variable name.

Intuitively, $X[\langle \text{pred exp} \rangle](x)$ represents a node of type X that satisfies the condition specified by $\langle \text{pred exp} \rangle$ and this node is denoted by the variable x . When

2. <http://lucene.apache.org>

```

<query> ::= <pattern> ':' (<link cond>)? ':' (<proximity cond>)? ':' (<return exp>)?
<pattern> ::= <vert axis> <node exp> ((' (<var> ')')? ('{' <pattern list> '})')?
<node exp> ::= <node> ('[' <pred exp> ']')? (<or op> <node exp>)*
<node> ::= <id> | '?'
<vert axis> ::= '/' | '/'
<pattern list> ::= <pattern> (<horz axis> <pattern>)*
<horz axis> ::= '->' | '=>' | '<=>'
<pred exp> ::= <pred term> (<and op> <pred exp>)*
<pred term> ::= <field> ( (<op> <value>) | (<IN op> '(' <value> (',' <value>)* ')') )
<link exp> ::= <link term> (<bool op> <link exp>)*
<link term> ::= (<var> <link type> <var>) | '(' <link exp> ')'
<link type> ::= '!<link_main_type> (<link_subtype>)?
<proximity exp> ::= <proximity term> (<bool op> <proximity exp>)*
<proximity term> ::= (<number>? ('[' | '{') <var> (<var>+
    (']' | '}') <number>?) | '(' <proximity exp> ')'
<return exp> ::= (DISTINCT)? <return term> (',' <return term>)*
<return term> ::= <var> '.' <field>
<op> ::= '=' | '<' | 'like' | 'is' | 'is not'
<and op> ::= 'AND'
<or op> ::= 'OR'
<in op> ::= 'IN'

```

Fig. 8. PTQL Grammar

the predicate expression is empty (resp. variable binding is empty) we can omit the square brackets (resp. parenthesis). As an example, $N[\text{tag}='P'](x)$ is a node expression describing a node labeled with N (noun) and tagged with P (i.e., a protein name) and this node is denoted by the variable x . A wildcard '?' denotes a node of any name. For instance, $?(y)$ represents a node of any name and this node is denoted by the variable y .

To describe a tree pattern, we use two types of axis. A vertical axis / (parent-child relationship) or // (ancestor-descendant relationship) describes the hierarchical order of nodes in the parse tree. A horizontal axis -> (immediate following) or => (following) describes the horizontal order of nodes³. Formally, a tree pattern is defined recursively as follows.

Definition 2: If e is a node expression then $/e$ and $//e$ are tree patterns. If e is a node expression and q_1, \dots, q_n are tree patterns then $/e\{q_1 \langle ha_1 \rangle \dots \langle ha_{n-1} \rangle q_n\}$ and $//e\{q_1 \langle ha_1 \rangle \dots \langle ha_{n-1} \rangle q_n\}$ are tree patterns, where ha_i can be ->, =>, or <=>.

A parse tree matches a pattern $/e$ (resp. $//e$) if one of the children (resp. descendants) of the root node satisfies the node expression e . A parse tree matches a pattern $/e\{q_1 \langle ha \rangle \dots \langle ha \rangle q_n\}$ (resp. $//e\{q_1 \langle ha_1 \rangle \dots \langle ha_{n-1} \rangle q_n\}$) if there is a node X with children Y_1, Y_2, \dots, Y_n of X such that (i) X is a child (resp. descendant) of the root node of the parse tree, (ii) X satisfies the node expression e , (iii) each tree T_i with X being the root node and the subtree of the parse tree rooted at Y_i being the only subtree of T_i matches the pattern q_i , and (iv) if $ha_i = '->'$ (resp. $ha_i = '=>'$) then Y_{i+1} immediately follows (resp. follows) Y_i . Y_{i+1} follows or precedes Y_i if $ha_i = '<=>'$. For

3. A node X is said to (immediately) follow a node Y in a parse tree if the right most leaf of X (immediately) follows the left most leaf of Y .

instance, the parse tree in Figure 6 matches the pattern $//S\{NP->/VP\}$ as (i) the node labeled with S is a descendant of the root node of the parse tree, (ii) the first noun phrase and the verb phrase are two children of that node, and (iii) the verb phrase immediately follows the noun phrase. This tree also matches the pattern $//S\{///?[tag='P']=>/VP\{V->///?[tag='P']\}$. However, it does not match the pattern $//S\{//NP->/V\}$ as the verb does not immediately follow either the first noun phrase (there is an adverb in between) or the second noun phrase. A link condition is defined as follows.

Definition 3: A link term is an expression of the form $x !\langle \text{link} \rangle y$, where x and y are variable names and $\langle \text{link} \rangle$ is a link name in the linkage. A link condition is a boolean expression of link terms.

For instance, $x !S y$ is a link term representing the fact that the node denoted by x connects to the node denoted by y through an S link. Similarly, $y !O z$ is a link term representing the fact that the node denoted by y connects to the node denoted by z through an O link. $x !S y$ AND $y !O z$ is a link expression whose meaning should be clear.

Definition 4: A proximity term is an expression of the form $m\{x_1 \dots x_k\}n$ or $m[x_1 \dots x_k]n$, where x_1, \dots, x_k are variable names and m, n are integers. A proximity condition is a boolean expression of proximity terms.

We use an example to illustrate the definition of proximity terms. $1\{x y\}2$ is a proximity term representing the fact that the nodes denoted by x, y are at least 1 node but not more than 2 nodes apart with respect to the sentence that contains words represented by x and y . In the case of $1[x y]2$, an additional constraint is imposed such that words represented by x have to appear before words represented by y in the sentence.

Definition 5: A *return expression* is a list of elements of the form $\langle \text{var} \rangle . \langle \text{attr} \rangle$ separated by ', ', possibly preceded by the keyword `DISTINCT`, where $\langle \text{var} \rangle$ is a variable name and $\langle \text{attr} \rangle$ is an attribute name.

As an example, `DISTINCT x.value, y.value, z.value` is a return expression that returns the distinct value attributes of nodes denoted by variables x, y, z in the tree pattern. We now define the syntax of PTQL queries.

Definition 6: A *PTQL query* is an expression of the form $\langle \text{pattern} \rangle : \langle \text{link cond} \rangle : \langle \text{proximity cond} \rangle : \langle \text{return exp} \rangle$ where $\langle \text{pattern} \rangle$ is a tree pattern, $\langle \text{link cond} \rangle$ is a link condition, $\langle \text{proximity cond} \rangle$ is a proximity condition and $\langle \text{return exp} \rangle$ is a return expression.

A parse tree matches a PTQL query if it matches the tree pattern of the query and the links between nodes satisfies the link condition of the query. The *return expression* of the query defines what information we want to return.

We illustrate the PTQL queries with examples, as shown in Table 1. For query Q1, the tree pattern `//?[tag='P'](i1)=>//V[value='regulates'](v)=>//?[tag='P'](i2)` represents that a protein name (denoted as $i1$) is followed by the verb “regulates” (denoted as v), which is followed by another protein name (denoted as $i2$), while the link condition `i1 !S v` and `v !O i2` specifies that $i1$ has to be the subject of the sentence and v corresponds to the verb “regulates”, and v and $i2$ has verb-object relation. This query returns the triplet (RAD53, regulates, DBF4). Query Q2 demonstrates how tree patterns can be used to identify protein mentions based on lexical clues using the word “protein” within a noun phrase. This query returns (DBF4) as a protein mention. Query Q3 utilizes the link `MX` to resolve coreferences, as the link `MX` connects a relative pronoun to its corresponding word. In this case, the relative pronoun “which” is connected to the word “RAD53” so that (RAD53) is returned.

3.3 Query evaluation

Our approach for the evaluation of PTQL queries involves the use of IR engine as well as RDBMS. The role of the IR engine in query is to select sentences based on the lexical features defined in PTQL queries, and only the subset of sentences retrieved by the IR engine are considered for the evaluation of the conditions specified in the PTQL queries by RDBMS. Unlike the filtering mechanism described in [17] that selects potentially relevant documents for extraction, our approach does not discard sentences that should otherwise be included for extraction. Using sample query Q1 as an example, the lexical features defined in the query imply that only sentences with at least one gene name together with the keyword “regulates” should be considered for

extraction. We summarize the process of the evaluation of PTQL queries as follows:

- 1) Translate the given PTQL query into a filtering query.
- 2) Use the filtering query to retrieve relevant documents D and the corresponding sentences S from the inverted index.
- 3) Translate the PTQL query into an SQL query and instantiate the query with document id $d \in D$ and sentence id $s \in S$.
- 4) Query PTDB using the SQL query generated in Step 3.
- 5) Return the results of the SQL query as the results of the PTQL query.

In step 2, the process of finding relevant sentences with respect to the given PTQL query requires the translation of the PTQL query into the corresponding filtering query. Here we define the syntax of the *keyword-based filtering queries*, which adopts the syntax of Lucene queries.

A *query term* t for a filtering query is a string that can be preceded by the required operator $+$, as well as the term $\langle \text{field} \rangle :$, where $\langle \text{field} \rangle$ is the name of a field. A *phrase* p is in the form $"t_1 \dots t_n"$, where t_1, \dots, t_n are query terms. p can be followed by a proximity operator in the form of $p \sim \langle \text{number} \rangle$. A parenthesis expression is composed of query terms and phrases, enclosed by parentheses, and it can be preceded by the required operator. A *keyword-based filtering query* is a list of query terms, phrases and parenthesis expressions. An PTQL query q is translated into a keyword-based filtering query using the following steps:

- 1) Generate query terms for each of the node expressions that are in the tree pattern of q .
- 2) Form phrases if consecutive node expressions are connected by “immediate following” horizontal axes (i.e. “ \rightarrow ”).
- 3) Form phrases followed by the proximity operator if the corresponding nodes are defined in the proximity condition of q .

The translation of a PTQL query q into a keyword-based filtering query involves the traversal of the parse tree of the PTQL query in preorder walk fashion. For each predicate term in the form of $\langle \text{field} \rangle = \langle \text{val} \rangle$ (resp. `IN (val1, ..., valk)`), we create the query term $+\langle \text{field} \rangle : \langle \text{val} \rangle$ (resp. $+\langle \text{field} \rangle : (\langle \text{val}_1 \rangle \dots \langle \text{val}_k \rangle)$)⁴. For example, the predicate term of a node expression `//N[tag='P' and value IN ('RAD53', 'DBF4')]` is `+sent: ("RAD53" "DBF4")+rep-sent:PROTNAME`. If the tree pattern of q contains $e_i \rightarrow e_j$, where e_i and e_j are node expressions with predicate term in the form of $\langle \text{field} \rangle = \langle \text{val} \rangle$, then the phrase $"\langle \text{val}_i \rangle \langle \text{val}_j \rangle"$ is formed for the query terms that represent e_i and e_j . If a proximity

4. $+\langle \text{field} \rangle : (\langle \text{val}_1 \rangle \dots \langle \text{val}_k \rangle)$ is the short form for $+(\langle \text{field} \rangle : \langle \text{val}_1 \rangle \dots \langle \text{field} \rangle : \langle \text{val}_k \rangle)$

<i>Relationship extraction (Q1):</i>	Extract protein-protein interactions with the pattern <subject>-<verb>-<object>, where the subject and object correspond to protein names (tag='P') and the verb corresponds to "regulates"
PTQL:	//S{/?[tag='P'](i1)=//V[value='regulates'](v)=//?[tag='P'](i2)}: i1 !S v and v !O i2 :: i1.value, v.value, i2.value
<i>Entity recognition (Q2):</i>	Find protein mentions that precede the word "protein" within a noun phrase so that the return values of g are treated as protein names
PTQL:	//NP{/?(g)->/N[value='protein']} :: g.value
<i>Coreference resolution (Q3):</i>	Find words that are referenced by the pronoun "which" so that the corresponding words are the return values of w
PTQL:	//S{/?(w)=//PRP[value='which'](p)} : w !MX p :: w.value

TABLE 1

Examples of PTQL queries. The constituent tree and linkage given in Figures 3 and 4 match any of these queries.

term in the proximity condition of q is in the form $[x_i x_j]_n$, then the phrase " $\langle val_i \rangle \langle val_j \rangle$ " $\sim n$ is formed.

We use query Q1 in Table 1 to illustrate the translation of PTQL queries into keyword-based filtering. The following is the keyword-based filtering query for Q1:

```
+sent:regulates +rep_sent:PROTNAME
```

The translation of PTQL queries into SQL queries in step 3 adopts the approach used in LPath [5], [13] to translate the hierarchical representation and horizontal relations of PTQL queries into nested SQL queries. We further extend the translation of PTQL link conditions into SQL. The details of the PTQL-to-SQL translation is presented in Appendices C and D. The translated SQL query is then instantiated with the sentences retrieved in step 2 by specifying the document (Doc_ID) and sentence IDs (Sent_CID) in the SQL queries. The corresponding SQL translation for query Q1 is as follows, in which table alias C and T refers to the table Constituents and Linkages that are used for storing the constituent trees and linkages of sentences, while table alias B refers to the table Bioentities that is used for storing entity mentions:

```
SELECT v2,v3,v4
FROM (SELECT id id4,...,v v4,id1,...,v1
      FROM C,(SELECT id id3,...,v v3,id1,...,v1
            FROM C,(SELECT id id2,...,v v2,id1,...,v1
                  FROM C,(SELECT pid pid1,...,v v1
                        FROM C WHERE d>2 AND C.n='S') C2
                  WHERE C.l>=C2.l1 AND C.r<=C2.r1
                        AND C.d>C2.d1 AND (B.t='P' AND
                        B.id=C.id)) C3
            WHERE C.l>=C3.l1 AND C.r<=C3.r1
            AND C.d>C3.d1 AND C.n='V'
            AND C.v='regulates'
            AND C.l>=C3.r2) C4
      WHERE C.l>=C4.l1 AND C.r<=C4.r1 AND
            C.d>C4.d1 AND (B.t='P' AND
            B.id=C.id) AND C.l>=C4.r3) T
WHERE ((T.id2,T.id3) IN (SELECT f_id,t_id
                        FROM L WHERE TYPE='S')) AND ((T.id3,T.id4)
                        IN (SELECT f_id,t_id FROM L WHERE TYPE='O'))
```

4 QUERY GENERATION

In this section, we demonstrate how PTQL queries for extraction differs from the traditional pipeline approach of information extraction through the extraction of

protein-protein and drug-protein (also known as drug-enzyme) interactions. Information extraction is typically performed by identifying a set of distinctive features that are common among the examples of the relationship of interest from the training data. These distinctive features are then utilized to form extraction patterns that are generic enough to be applied to unseen data for relationship extraction. We called this approach as the *training set driven query generation*, and we demonstrate its capability in extracting any kinds of protein-protein interactions. However, compiling a corpus of annotated training data is known to be time-consuming and labor-intensive. In the absence of training data, our alternative approach is to generate extraction queries from keyword-based queries using a pseudo-relevance feedback methodology. We called this alternative approach as the *pseudo-relevance feedback driven query generation* that is capable of extracting specific kind of interactions. In this case, we show the performance of the approach with the extraction of drug-protein metabolic relations.

4.1 Training set driven query generation

We illustrate our approach with an application of protein-protein interaction extraction using a set of syntactic patterns that are expressed in PTQL queries. To generate a set of patterns for information extraction, the *annotator* component is applied to automatically annotate an unlabeled document collection with information drawn from a problem-specific database. This step necessitates a method for precise recognition and normalization of protein mentions. From this labeled data, the *pattern generator* identifies relevant phrases referring to interactions in order to generate patterns. These initial patterns are then used to compute consensus patterns through the *pattern generalization* component for protein-protein interactions (PPI)s. PTQL queries are then formed by the *query generator* to perform extraction from the parse tree database.

As many sentences might contain coincidental mentions of proteins and not describe an interaction ("We study the proteins A, B, and C."), we reduce these initial candidate evidence by a number of refinement steps. As a first step, we search for words typically referring to protein-protein interactions ("binds", "association", "mediated"). We currently use a set of 123 verbs, 126

1	P	interacts	with	the	P
2	P	binds	to		P
3	P	bound	to		P
	P	{i-verb}	{preposition}	{determiner}?	P

{determiner} := a, an, the, these, this, those, ...;
 {preposition} := between, for, to, with, ...;
 {i-verb} := binds, bound, interacts, interacted, ...;

Fig. 9. Multiple initial patterns (1–3) lead to the same general pattern, after words have been replaced with concepts (in curly brackets; for example, protein names are replaced with ‘P’).

nouns, and 8 adjectives, plus corresponding word forms. Such words have to appear between the two proteins under consideration, or precede/follow them in a short distance, which is parameterizable. We then reduce the full sentence to the snippet that likely conveys the information about an interaction; therefore, we may extract the shortest snippet that contains both proteins and an interaction-indicating word, or include additional words from the left and right of this snippet. Each snippet found is considered as a *relevant phrase*; it could be directly used to find similar (parts of) sentences in the test dataset. The more the snippet extends to the left and right, the more precise the phrase will be. Shorter snippets, on the other hand, will typically increase the recall when using that phrase.

To increase the recall of the initial patterns, we generalize the patterns by substituting tokens belonging to certain word categories with a higher-level concept, such as gene/protein names, interaction verbs. It shows that many natural language sentences that describe certain events, like protein-protein interactions, exhibit a certain ‘behavior’ that makes them similar to other such sentences. Consider Figure 9, where the words ‘binds’, ‘interacts’, etc. can be generalized into *i-verb*, indicating that the words belong to the higher-level concept *interaction verbs*.

The generalized patterns are then translated into PTQL queries for extraction. For instance, the syntactic pattern $\langle P \rangle \{i\text{-verb}\} \{preposition\} \{determiner\}?\langle P \rangle$, where $\langle P \rangle$ corresponds to the matching of any protein names, is translated into:

```
//S{/[tag='P'](i1)->/V[tag='I']->/IN[value
  in {'with','for',...}(w1)=>/[tag='P'](i2)}
  :: 1[w1 i2]2 : i1.value, i2.value
```

The category $\{preposition\}$ is represented by its set of predefined instances of prepositions in the PTQL query. The term $\{determiner\}?$ indicates that a determiner is an optional match, and it is represented by the proximity condition $1[w1 i2]2$ so that at most one word can be in between the words represented by the variables $w1$ and $i2$. Notice that no link condition is used in the translation from syntactic patterns to PTQL queries.

4.2 Pseudo-relevance feedback driven query generation

To make up for the lack of training data for the relationship of interest, we offer an alternative approach that is inspired by pseudo-relevance feedback in information retrieval (IR). The idea is to automatically generate PTQL queries by considering the constituent trees of the top- k sentences retrieved with a boolean keyword-based query. The common grammatical patterns among the constituent trees of relevant sentences are utilized to form extraction patterns. Interaction extraction is then performed by using the PTQL queries translated from the generated extraction patterns.

A boolean keyword-based query q is composed of query terms $t_1 \dots t_n$, where a query term t_i can be a keyword, or an identifier for an entity type, such as `PROTNAME` that represents any matches of protein names. With q , a ranked list of sentences S is retrieved and the constituent trees of the top- k sentences of S (denoted as S_k) are retrieved from PTDB. To find common grammatical patterns among the constituent trees of S_k , string encodings are generated for each of the sentence in S_k . A 0 -th level string encoding records the labels of the lowest common ancestor lca of the query terms and the query terms themselves in a pre-order tree traverse order. A m -th level string encoding is defined as the string encoding that includes at most m descendants of lca on each of the paths connecting lca and a query term t_i . For instance, suppose $q=PROTNAME$ and `regulates` and `PROTNAME`, then the string

```
//S/NP/kw1://S/VP/kw2://S/VP//kw3
```

is a 1st level string encoding for the constituent tree in Figure 4, in which `kw1`, `kw3` represent `PROTNAME` and `kw2` represents the keyword “`regulates`”, and S is the the lowest common ancestor node for the three terms. A m -th level string encoding has a one-to-one translation to a PTQL query. Using the above string encoding as an example, the corresponding PTQL query is

```
//S{/NP{/[tag='P']}=>/VP{/[value=
  'regulates']=>/[tag='P']}}}
```

With $m=0$, the linguistic patterns in the relevant sentences are maximally generalized, potentially leading to a high recall with possible compromise to precision. With the increasing value of m , the patterns become more specific, potentially increase precision with possible compromise to recall. By forming the m -th level string encodings, we can identify the similarity of the retrieved sentences based on their grammatical structures. We define that two sentences are *grammatically similar* if they have the same m -th level string encoding. Grammatically similar sentences are grouped together to form a cluster. A PTQL query is then generated for each of the clusters of string encodings. The steps of generating PTQL queries can be outlined as follows: Let C_m be a set of clusters with m -th level string encodings. Given a boolean keyword-based query q and

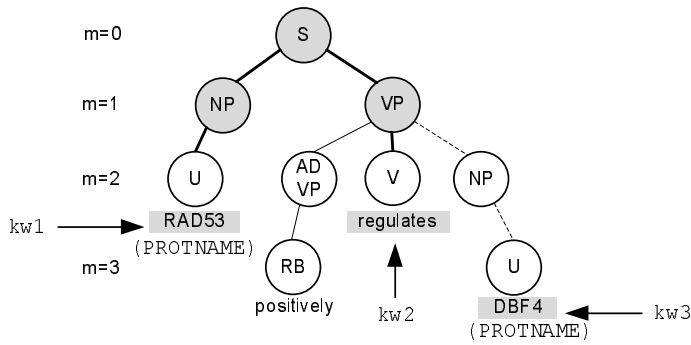


Fig. 10. An illustration of the m -th level string encoding $//S/NP/kw1://S/VP/kw2://S/VP//kw3$, where $m=1$, for the constituent tree of “RAD53 positively regulates DBF4” with respect to the query `PROTNAME` and `regulates` and `PROTNAME`. The shaded nodes in the constituent tree are used for the string encoding, and the node with label `s` is the the lowest common ancestor node for the three query terms.

parameter k ,

- 1) Retrieve sentences using q from the inverted index and retrieve the constituent trees of the top- k sentences S_k from PTDB.
- 2) For each sentence in S_k extract the subtree that is rooted at the lca of all the query terms t_1, \dots, t_n with the query terms as leaf nodes from the constituent tree.
- 3) Generate m -th level string encodings for each of the subtrees.
- 4) Sentences that are grammatically similar based on their m -th level string encodings are grouped together to form clusters of common grammatical patterns C_m .
- 5) A PTQL query is generated for each common grammatical pattern C_m .

Interactions are extracted through the evaluation of the generated PTQL queries.

5 RESULTS

We first illustrate the performance of our approach in terms of query evaluation and the time savings achieved through incremental extraction. Then we evaluate the extraction performance for our two approaches in query generation.

5.1 Time performance for PTQL

We performed experiments in finding the time performance of the evaluation of PTQL queries, as well as experiments to illustrate the amount of time saved in the event of change of an extraction goal and deployment of an improved module. All experiments were performed using a 2.2-GHz Intel Xeon QuadCore CPU running in Red Hat Linux. Only a single process was used to perform the experiments. The parse tree database is stored as a relational database managed by MySQL.




Fig. 11. Time performance in seconds for PTQL evaluation for a set of 25 queries that involve the extraction of drug-metabolic-gene relations.

5.1.1 Query evaluation

Using a large corpus of 17 million Medline abstracts, the intermediate processing data was stored in our parse tree database occupying about 1.5 TB of disk space. A set of 25 PTQL queries that involves the extraction of drug-gene metabolic relations was used to evaluate the time performance for PTQL evaluation. Specifically, given a drug, the goal is to find which genes are involved in the metabolic relations with the drug. In our experiments, we specified a single drug (“1-drug”), a set of 5 drugs (“5-drugs”) and a set of 10 drugs (“10-drugs”) in each of the 25 PTQL queries. Each query was evaluated with 5 different sets of drugs and repeated for 5 runs. Figure 11 shows the average duration over different sets of drugs and runs for the PTQL evaluation. This figure shows that the query evaluation can be completed in the range of a second to 50 seconds for all answers to be returned. Queries specified with a larger set of drugs require a longer time to complete the evaluation. The time performance indicates that our proposed framework is acceptable for real-time IE.

5.1.2 Incremental extraction

We performed experiments to show the time savings for incremental extraction. Incremental extraction is applied when there is a change in extraction goals or deployment of an improved processing module. The scenario behind our experiment is that the initial goal is to perform extraction for drug information from a text collection. The extraction goal is then changed into the extraction of drug-protein relations that requires the deployment of a gene named entity recognizing to identify gene mentions in the text collection.

To illustrate the amount of time savings, a collection of 13K Medline abstracts was initially processed with the Link Grammar parser and a dictionary-based tagger for drug names. This process took about 62.38 hours. We then deployed a statistical-based tagger for gene names to process the corpus. With the pipeline approach, the whole process had to be started from scratch by running

the Link Grammar parser, the drug name tagger and the newly deployed gene name tagger. This took another 64.8 hours to complete. With our approach, the intermediate processing data produced by the Link Grammar parser and the drug name tagger were populated into the parse tree database. The gene name tagger was then deployed to process the corpus and SQL insert statements were issued to update the parse tree database. This process took only 6.71 hours to complete. This experiment showed a tremendous decrease of 89.64% when a new module is deployed for text processing as compared to the pipeline approach.

5.2 Extraction performance for PTQL

We used two datasets to evaluate the performance of PTQL. For the evaluation of our training set driven query generation, we used the BioCreative 2 IPS test data [18] as a benchmark for relationship, in this case the extraction of protein-protein interactions. This data set consists of 358 full-text articles, which we transformed into 98,209 sentences. This set was reduced by us to include only sentences that contain at least one protein, resulting in 71,631 sentences. Another dataset of 13015 Medline abstracts from [19] that focus on drug-protein relations is used to evaluate the pseudo-relevance feedback driven query generation.

5.2.1 Training set driven query generation

Using the training set driven query generation method, we used the training data from the BioCreative 2 IPS corpus to generate extraction patterns. The task in the BioCreative 2 IPS benchmark is to find protein-protein interactions for which a text provides evidence for a physical interaction between the proteins. In addition, all proteins have to be mapped to corresponding identifiers in the UniProt database. We describe our approach for this additional task in [20]. The generated patterns were then translated into PTQL queries, which were utilized to perform extraction from the testing data. The training set driven query generation method generated 11,208 extraction patterns. Our approach of using PTQL queries to express the generated patterns achieves a mean precision and recall of 83.6% and 58.6% with 64.2% as the *f*-measure. The extraction results are summarized in Table 2. Our results also show far better performance than the previously top-performing PPI systems [21], [22], [23]. To give a fair comparison, we performed another experiment that utilized the 74 manually curated patterns reported in [24] on their OpenDMAP system, but using the same gene normalization (EMN) and named entity recognition (NER) from [20]. We observed that our approach still achieves significantly better results. The significant improvement over previous methods is largely contributed by the large number of generated extraction patterns that were generated by our training set driven query generation approach. Each of the generated patterns has a high precision but low recall,

TABLE 2

Performance of various approaches on the BioCreative 2 IPS test data (version: “SwissProt only”). Mean precision, recall, and *f*-measure in %. The systems marked * are basically the same; for one, we replaced their NER/EMN with our own to get results that are better comparable.

System	P	R	F
PTQL	83.6	58.6	64.2
* Reported by [21]	39.1	29.7	28.5
74 manually created patterns [24], with NER/EMN from [20]	59.7	37.9	41.8
Reported by [22]	37.0	32.7	30.4
Reported by [23]	25.2	23.3	24.2

but the combined results of the 11,208 extraction queries contributes to a high overall recall.

5.2.2 Pseudo-relevance feedback query generation

We further evaluate the extraction performance for PTQL using the pseudo-relevance feedback driven query generation method. Our goal here is to extract drug-protein metabolic relations without the use of training data to generate queries. Specifically we created boolean keyword-based queries in the form of DRUGNAME and <metabolic-word> and PROTNAME, where the identifiers DRUGNAME and PROTNAME correspond to matching of any mentions of drug and protein names, and <metabolic-word> is a class of words to indicate drug-protein metabolic relations that include the words “metabolized”, “metabolize”, “metabolizes”, “metabolised”, “metabolise”, “metabolises” and “metabolism”. For each of the keyword-based queries, the top *k*% of the retrieved sentences were used to generate clusters. A cluster is translated into a PTQL query if it has at least *n* members. The generated PTQL queries are then applied to perform extraction, and the query results of the PTQL queries are the extracted relations.

The performance of the query generation using our pseudo-relevance feedback is compared with the cooccurrences method, which considers a drug-protein metabolic relation when drug and protein names together appear with one of the metabolic words in a sentence. To perform the comparison, we created a gold standard by analyzing all of the interactions that were extracted using the cooccurrences method. This results in a collection of 1059 drug-protein metabolic relations⁵ out of the 13,015 abstracts. Table 3 shows the performance of the query generation method compared to the cooccurrences method. While the *f*-measure shows that the query generation method lags behind due to the high recall achieved by cooccurrences, it is important to notice that precision is significantly higher than its cooccurrences counterpart. The results also shows the expected tradeoff between precision and recall when *m* varies. With increasing *m*, which is the maximum number of

5. We omit possible cross-sentence relations in the creation of gold standard.

TABLE 3

Precision, recall and f-measure in % for drug-protein metabolic relations between the cooccurrences method and the pseudo-relevance feedback driven query generation method (QueryGen). m is the maximum number of descendants to include in the m -th level string encodings, using the top 60% of the retrieved sentences to form clusters with at least 3 members.

System	Precision	Recall	F-measure
Cooccurrences	39.9	100.0	57.0
QueryGen ($m=1$)	53.5	64.3	58.4
QueryGen ($m=2$)	68.0	42.8	52.5
QueryGen ($m=3$)	83.3	31.0	45.2
QueryGen ($m=4$)	85.0	26.0	39.8

TABLE 4

Comparison of cooccurrence and pseudo-relevance feedback on the extraction of drug-protein metabolic relations. Number of true positives (TP), precision (P), recall (R) and f-measure (F) for each kind of gene-drug metabolic relations between the cooccurrences method and the pseudo-relevance feedback driven query generation method (PTQL). m is the maximum number of descendants to include in the m -th level string encodings, using the top 60% of the retrieved sentences to form clusters with at least 3 members.

	Cooccurrences				PTQL ($m=1$)			
	TP	P	R	F	TP	P	R	F
<i>metabolized</i>	344	55.0	100.0	71.0	243	82.2	68.6	74.8
<i>metabolises</i>	9	45.0	100.0	62.1	–	–	–	–
<i>metabolised</i>	82	51.9	100.0	68.3	43	79.6	52.4	63.2
<i>metabolize</i>	38	40.9	100.0	58.0	13	52.0	34.2	41.3
<i>metabolism</i>	546	32.0	100.0	48.4	355	41.0	65.0	50.3
<i>metabolizes</i>	47	72.3	100.0	83.9	26	83.9	55.3	66.7

descendants of LCA to include in the m -th level string encodings, precision gains at the expense of recall. This is intuitive as the string encodings become more specific when more descendants are included. Table 4 shows the extraction performance of drug-protein metabolic relations with different keywords using our pseudo-relevance feedback driven query generation method. Extraction with the keyword “*metabolises*” does not yield any results with our query generation method, due to the small number of sentences retrieved by the initial keyword-based queries.

6 RELATED WORK

Information extraction has been an active research area over the years. The main focus has been on improving the accuracy of the extraction systems, and IE has been seen as an one-time execution process. A recent special issue of SIGMOD Record [25] highlights the need to develop IE frameworks that manage extraction tasks as long-time running processes.

6.1 Traditional IE approaches

Our proposed framework has a fundamental difference from the existing extraction systems in terms of the usage

of RDBMS. Popular IE frameworks such as UIMA [1] and GATE [2] provide the ability of efficient integration of various NLP components for IE. Such frameworks are file-based and they do not store the intermediate processing output of various components. Typical extraction systems such as QXtract [17] and Snowball [26] utilize RDBMS to store and query the extracted facts. Recent work on IE management systems rely on RDBMS for optimization of the execution of IE tasks. For systems such as Cimple [27] and SystemT [28], operations such as joins in RDBMS are performed over extracted facts that are stored in various database tables. However, in the event of a deployment of an improved component or a change of extraction goals, all components have to be reprocessed from scratch, which is computationally expensive as shown in our experiments. In our approach, only the new module has to be processed for the text collection. The intermediate processing data is then inserted into the parse tree database so that both the new and existing processing data can be utilized for extraction.

To address the high computational cost associated with extraction, document filtering is a common approach in which only the promising documents are considered for extraction [17], [29], [30]. These promising documents are selected based on a classifier that is trained for determining documents that are relevant for extraction. Such an approach can potentially miss out documents that should have been used for extraction. In our filtering approach, sentences are selected solely based on the lexical clues that are provided in a PTQL query. This filtering process utilizes the efficiency of IR engines so that a complete scan of the parse tree database is unnecessary without sacrificing any sentences that should have been used for extraction.

6.2 Rule-based IE approaches

Rule-based IE approaches have been proposed in [31], [32], [33], [34]. The AQL query language proposed in the Avatar system [34] is capable of performing extraction for a variety of IE tasks that includes matching with regular expression. Unlike PTQL, the language does not support the use of parse trees, which can be useful in IE tasks such as relationship extraction. Systems such as DIAL [32], TLM [33], KnowItNow [35] focus on relationship extraction based on their own query languages. However, these languages only support querying of data from shallow parsing, and they do not have the ability of extracting from rich grammatical structures such as parse trees. On the other hand, declarative languages are used in systems such as Cimple [27] and SystemT [28]. Operations such as joins in RDBMS are performed over extracted facts that are stored in various database tables. Rules are then applied to integrate different types of extracted facts. However, these rules are not capable of querying parse trees.

Our work is closest to an IE management system MEDIE [31] stores parse trees in a database and allows

extraction over parse trees with a query language. The XML-like query language proposed is based on another kind of dependency grammar called head-driven phrase structure (HPSG). Unlike PTQL, link types cannot be expressed with this query language. In addition, MEDIE only provides simple query generation component that translates subject-verb-object extraction queries into its own query languages. This limits the utility of the system, as there can be a learning curve to learn the query language before users can perform their own extraction.

6.3 Machine learning approaches for IE

Our proposed approach provides mechanisms to generate extraction queries from both labeled and unlabeled data. Query generation is critical so that casual users can specify their information needs without learning the query language. Learning extraction patterns from training data has been proposed previously [17], [26], [36], [37]. However, training data is not always readily available and annotation of training data is both labor-intensive and time-consuming. To compensate the need of training data for extraction, [38], [39] proposed a new paradigm of automated discovery of relations through self-supervised learning. The core idea is to identify dependency structures from sentences that are retrieved by a user's query. The identified grammatical structures are generalized to form extraction patterns. However, the process of finding patterns can be expensive as any subtrees in the dependency structures can be considered as patterns. To reduce the complexity, only structures with verbs as predicates and a restricted number of nodes are considered in the process of identifying patterns. The techniques in [38], [39] requires certain predefined linguistic heuristics for the generation of extraction patterns. Our query generation method does not make any assumptions for linguistic heuristics.

7 DISCUSSION AND FUTURE WORK

In this section, we discuss the main contributions of our work as well as their limitations.

- *Extraction framework.* Existing extraction frameworks do not provide the capabilities of managing intermediate processed data such as parse trees and semantic information. This leads to the need of reprocessing of the entire text collection, which can be computationally expensive. On the other hand, by storing the intermediate processed data as in our novel framework, introducing new knowledge can be issued with simple SQL insert statements on top of the processed data. With the use of parse trees, our framework is most suitable for performing extraction on text corpus written in natural sentences such as the biomedical literature. In the case when the parser fails to generate parse tree for a sentence, our system generates a "replacement parse tree" that has the node `STN` as the root with the words in the sentence as the children of the root node. This allows PTQL queries to be applied to sentences that are incomplete or casually-written, which

can appear frequently in web documents. Features such as horizontal axis and proximity conditions can be most useful for performing extraction on replacement parse trees.

- *Parse tree query language.* One of the main contributions of our work is PTQL that enables information extraction over parse trees. While our current focus is per-sentence extraction, it is important to notice that the query language itself is capable of defining patterns across multiple sentences. By storing documents in the form of parse trees, in which the node `DOC` is represented as the root of the document and the sentences represented by the nodes `STN` as the descendants. As shown in the sample queries illustrated in Table 1, PTQL has the ability to perform a variety of information extraction tasks by taking advantage of parse trees unlike other query languages. Currently PTQL lacks the support of common features such as regular expression as frequently used by entity extraction task. PTQL also does not provide the ability to compute statistics across multiple extraction such as taking redundancy into account for boosting the confidence of an extracted fact.

- *Query generation* An important component of our extraction framework is on the automated generation of PTQL queries. Our training set driven query generation component is capable of generating precise extraction queries. To handle recall, the strategy is to generate a large number of such precise extraction queries. While training data is available for certain extraction goals such as the extraction of protein-protein interactions, it is not the case for other kinds of extraction. Our pseudo-relevance feedback query generation mechanism allows generation of PTQL queries without training data. As indicated by the precision listed in Table 4, the performance of the query generation method can be sensitive to the choice of keywords. However, our results indicate that our method can significantly improve the precision over the precision achieved by using cooccurrences.

As future work, we will extend the support of other parsers by providing wrappers of other dependency parsers and scheme, such as Pro3Gres and the Stanford Dependency scheme, so that they can be stored in PTDB and queried using PTQL. We will expand the capabilities of PTQL, such as the support of regular expression and the utilization of redundancy to compute confidence of the extracted information.

REFERENCES

- [1] D. Ferrucci and A. Lally, "UIMA: an architectural approach to unstructured information processing in the corporate research environment," *Natural Language Engineering*, vol. 10, no. 3-4, pp. 327-348, 2004.
- [2] H. Cunningham, D. Maynard, K. Bontcheva, and V. Tablan, "GATE: A framework and graphical development environment for robust NLP tools and applications," in *Proc. of ACL*, 2002.
- [3] A. Doan, J. F. Naughton, R. Ramakrishnan, A. Baid, X. Chai, F. Chen, T. Chen, E. Chu, P. DeRose, B. Gao, C. Gokhale, J. Huang, W. Shen, and B.-Q. Vuong, "Information extraction challenges in managing unstructured data," *SIGMOD Rec.*, vol. 37, no. 4, pp. 14-20, 2008.

- [4] D. Grinberg, J. Lafferty, and D. Sleator, "A robust parsing algorithm for Link Grammars," Pittsburgh, PA, Tech. Rep. CMU-CS-TR-95-125, 1995.
- [5] S. Bird, Y. Chen, and et. al, "Designing and evaluating an XPath dialect for linguistic queries," in *ICDE '06*, 2006.
- [6] S. Sarawagi, "Information extraction," *Foundations and Trends in Databases*, vol. 1, no. 3, pp. 261–377, 2008.
- [7] D. D. Sleator and D. Temperley, "Parsing English with a Link Grammar," in *Third Intl. Workshop on Parsing Technologies*, 1993.
- [8] R. Leaman and G. Gonzalez, "Banner: An executable survey of advances in biomedical named entity recognition," in *Pacific Symposium on Biocomputing 13*, 2008, pp. 652–663.
- [9] A. R. Aronson, "Effective mapping of biomedical text to the umls metathesaurus: the metemap program." in *Proceedings of the AMIA Symposium*. American Medical Informatics Association, 2001, p. 17.
- [10] M. J. Cafarella and O. Etzioni, "A search engine for natural language applications," in *WWW'05*, 2005.
- [11] T. Cheng and K. C.-C. Chang, "Entity search engine: Towards agile best-effort information integration over the web," in *CIDR*, 2007.
- [12] H. Bast and I. Weber, "The CompleteSearch Engine: Interactive, efficient, and towards IR& DB integration," in *CIDR*, 2007, pp. 88–95.
- [13] S. Bird, Y. Chen, S. B. Davidson, H. Lee, and Y. Zheng, "Extending xpath to support linguistic queries," in *Workshop on Programming Language Technologies for XML (PLAN-X)*, 2005.
- [14] J. Clark and S. DeRose, "XML Path language (XPath)," November 1999, <http://www.w3.org/TR/xpath>.
- [15] "XQuery 1.0: An XML query language," June 2001, <http://www.w3.org/XML/Query>.
- [16] C. Lai, "A formal framework for linguistic tree query," Department of Computer Science and Software Engineering, University of Melbourne, Master's thesis, 2005.
- [17] E. Agichtein and L. Gravano, "Querying text databases for efficient information extraction," in *ICDE*, 2003, pp. 113–124.
- [18] M. Krallinger, F. Leitner, and A. Valencia, "Assessment of the second Biocreative PPI task: Automatic extraction of protein-protein interactions," in *Proc 2nd BioCreative Challenge Evaluation*, 2007, pp. 41–54.
- [19] J. T. Chang and R. B. Altman, "Extracting and characterizing gene-drug relationships from the literature," *Pharmacogen.*, vol. 14, no. 9, pp. 577–586, Sep 2004.
- [20] J. Hakenberg, C. Plake, and et. al, "Inter-species normalization of gene mentions with GNAT," in *European Conference on Computational Biology (ECCB'08)*, 2008.
- [21] W. Baumgartner, Z. Lu, H. Johnson, J. Caporaso, and et al. J. Paquette, "An integrated approach to concept recognition in biomedical text," in *Proc. of the Second BioCreative Challenge Evaluation Workshop*, 2006.
- [22] M. Huang, S. Ding, H. Wang, and X. Zhu, "Mining physical protein-protein interactions by exploiting abundant features," in *Proc 2nd BioCreative Challenge Evaluation Workshop*, 2007, pp. 237–245.
- [23] J. Hakenberg, C. Plake, L. Royer, H. Strobelt, U. Leser, and M. Schroeder, "Gene mention normalization and interaction extraction with context models and sentence motifs," *Genome Biology*, vol. 9, no. Suppl 2, p. S14, 2008.
- [24] L. Hunter, Z. Lu, J. Firby, W. Baumgartner, H. Johnson, P. Ogren, and K. B. Cohen, "OpenDMAP: An open source, ontology-driven concept analysis engine, with applications to capturing knowledge regarding protein transport, protein interactions and cell-type-specific gene expression," *BMC Bioinformatics*, vol. 9, no. 1, p. 78, 2008.
- [25] A. Doan, L. Gravano, R. Ramakrishnan, and S. Vaithyanathan, "Introduction to the Special Issue on Managing Information Extraction," *SIGMOD Record*, vol. 37, no. 4, p. 5, 2008.
- [26] E. Agichtein and L. Gravano, "Snowball: Extracting relations from large plain-text collections," in *Proceedings of the fifth ACM conference on Digital libraries*. ACM New York, NY, USA, 2000, pp. 85–94.
- [27] A. Doan, J. Naughton, R. Ramakrishnan, A. Baid, X. Chai, F. Chen, T. Chen, E. Chu, P. DeRose, B. Gao et al., "Information extraction challenges in managing unstructured data," *ACM SIGMOD Record*, vol. 37, no. 4, pp. 14–20, 2009.
- [28] R. Krishnamurthy, Y. Li, S. Raghavan, F. Reiss, S. Vaithyanathan, and H. Zhu, "SystemT: a system for declarative information extraction," *ACM SIGMOD Record*, vol. 37, no. 4, pp. 7–13, 2009.
- [29] P. G. Ipeirotis, E. Agichtein, P. Jain, and L. Gravano, "Towards a query optimizer for text-centric tasks," *ACM Trans. DB Syst.*, vol. 32, no. 4, p. 21, 2007.
- [30] A. Jain, A. Doan, and L. Gravano, "Optimizing SQL queries over text databases," in *ICDE*, 2008.
- [31] Y. Miyao, T. Ohta, and et. al, "Semantic retrieval for the accurate identification of relational concepts in massive textbases," in *Proc. of ACL '06*, 2006, pp. 1017–1024.
- [32] R. Feldman, Y. Regev, E. Hurvitz, and M. Finkelstein-Landau, "Mining the biomedical literature using semantic analysis and natural language processing techniques," 2003.
- [33] J. D. Martin, "Fast and furious text mining," *IEEE Data Eng. Bull.*, vol. 28, no. 4, pp. 11–20, 2005.
- [34] F. Reiss, S. Raghavan, R. Krishnamurthy, H. Zhu, and S. Vaithyanathan, "An algebraic approach to rule-based information extraction," *Proc. ICDE08*, 2008.
- [35] M. Cafarella, D. Downey, S. Soderland, and O. Etzioni, "Knowitnow: Fast, scalable information extraction from the web," in *HLT-EMNLP-05*, 2005, pp. 563–570.
- [36] F. Suchanek, G. Ifrim, and G. Weikum, "LEILA: Learning to extract information by linguistic analysis," in *Proceedings of the ACL-06 Workshop on Ontology Learning and Population*, 2006, pp. 18–25.
- [37] F. Peng and A. Mccallum, "Accurate information extraction from research papers using conditional random fields," in *In HLT-NAACL*, 2004, pp. 329–336.
- [38] S. Sekine, "On-demand information extraction," in *Proceedings of the COLING/ACL on Main conference poster sessions*, 2006, pp. 731–738.
- [39] M. Banko, M. J. Cafarella, S. Soderland, M. Broadhead, and O. Etzioni, "Open information extraction from the web," in *Proc. of IJCAI*, 2007.

Luis Tari received the PhD degree in computer science from Arizona State University. He is currently a postdoc research fellow in text mining at Hoffmann-La Roche. His research interests include information extraction, information retrieval and bioinformatics.

Phan Huy Tu received the PhD degree in computer science from New Mexico State University. He is currently at Microsoft. His research interests include information extraction, information retrieval, reasoning, planning, knowledge representation and logic programming.

Jörg Hakenberg received the PhD degree in computer science from Humboldt-Universität zu Berlin. He is currently a research associate at Arizona State University. His research interests include text mining and natural language processing for biomedical applications.

Yi Chen received the PhD degree in computer science from University of Pennsylvania. She is currently an assistant professor of computer science at Arizona State University. Her research interests include keyword search on structured and semi-structured data, workflows, social network, information extraction and integration, and data streams.

Tran Cao Son received the PhD degree in computer science from New Mexico State University. He is currently an associate professor of computer science at New Mexico State University. His research interests include planning, knowledge representation and reasoning, logic programming, autonomous agents and web agents.

Graciela Gonzalez received the PhD degree in computer science from the University of Texas at El Paso. She is currently an assistant professor of biomedical informatics at Arizona State University. Her research interests include biomedical informatics, text mining, multimedia databases and human-computer interaction.

Chitta Baral received the PhD degree in computer science from University of Maryland. He is currently a professor of computer science at Arizona State University. His research interests include knowledge representation and reasoning, logic programming, bioinformatics, autonomous agents and reasoning about actions.