




(86) Date de dépôt PCT/PCT Filing Date: 2006/11/16
(87) Date publication PCT/PCT Publication Date: 2007/05/24
(45) Date de délivrance/Issue Date: 2016/05/24
(85) Entrée phase nationale/National Entry: 2009/05/08
(86) N° demande PCT/PCT Application No.: US 2006/044516
(87) N° publication PCT/PCT Publication No.: 2007/059287
(30) Priorité/Priority: 2005/11/16 (US60/737,446)

(51) Cl.Int./Int.Cl. *G06F 17/30* (2006.01)
(72) Inventeurs/Inventors:
MARCHISIO, GIOVANNI B., US;
NAVDEEP, DHILLON S., US;
TUSK, CARSTEN, US;
KOPERSKI, KRZYSZTOF, US;
LIANG, JISHENG, US;
NGUYEN, THIEN, US;
BROWN, MATTHEW E., US
(73) Propriétaire/Owner:
VCVC III LLC, US
(74) Agent: GOWLING WLG (CANADA) LLP

(54) Titre : EXTENSION DE RECHERCHE PAR MOT CLE A DES DONNEES ANNOTEES SYNTAXIQUEMENT ET SEMANTIQUEMENT
(54) Title: EXTENDING KEYWORD SEARCHING TO SYNTACTICALLY AND SEMANTICALLY ANNOTATED DATA



(57) Abrégé/Abstract:

Methods and systems for extending keyword searching techniques to syntactically and semantically annotated data are provided. Example embodiments provide a Syntactic Query Engine ("SQE") that parses, indexes, and stores a data set as an enhanced document index with document terms as well as information pertaining to the grammatical roles of the terms and ontological and other semantic information. In one embodiment, the enhanced document index is a form of term-clause index, that indexes terms and syntactic and semantic annotations at the clause level. The enhanced document index permits the use of a traditional keyword search engine to process relationship queries as well as to process standard document level keyword searches. This abstract is provided to comply with rules requiring an abstract, and it is submitted with the intention that it will not be used to interpret or limit the scope or meaning of the claims.



(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 May 2007 (24.05.2007)

PCT

(10) International Publication Number
WO 2007/059287 A1

(51) International Patent Classification:
G06F 17/30 (2006.01)

(21) International Application Number:

PCT/US2006/044516

(22) International Filing Date:

16 November 2006 (16.11.2006)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

60/737,446 16 November 2005 (16.11.2005) US

(71) Applicant (for all designated States except US): **IN-SIGHTFUL CORPORATION** [US/US]; 1700 Westlake Avenue North, Suite 500, Seattle, WA 98109-3044 (US).

(72) Inventor; and

(75) Inventor/Applicant (for US only): **MARCHISIO, Giovanni, B.** [US/US]; 17021 NE 37th Place, Bellevue, WA 98008 (US).

(74) Agents: **BIERMAN, Ellen, M.** et al.; SEED INTELLECTUAL PROPERTY LAW GROUP PLLC, Suite 5400, 701 Fifth Avenue, Seattle, WA 98104-7064 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.


(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: EXTENDING KEYWORD SEARCHING TO SYNTACTICALLY AND SEMANTICALLY ANNOTATED DATA



(57) Abstract: Methods and systems for extending keyword searching techniques to syntactically and semantically annotated data are provided. Example embodiments provide a Syntactic Query Engine ("SQE") that parses, indexes, and stores a data set as an enhanced document index with document terms as well as information pertaining to the grammatical roles of the terms and ontological and other semantic information. In one embodiment, the enhanced document index is a form of term-clause index, that indexes terms and syntactic and semantic annotations at the clause level. The enhanced document index permits the use of a traditional keyword search engine to process relationship queries as well as to process standard document level keyword searches. This abstract is provided to comply with rules requiring an abstract, and it is submitted with the intention that it will not be used to interpret or limit the scope or meaning of the claims.

WO 2007/059287 A1

EXTENDING KEYWORD SEARCHING TO SYNTACTICALLY AND SEMANTICALLY ANNOTATED DATA

TECHNICAL FIELD

The present invention relates to methods, systems, and techniques
5 for searching for information in a data set, and, in particular, to enhanced methods,
systems, and techniques for syntactically indexing and performing syntactic
searching of data sets using relationship queries to achieve greater search result
accuracy.

BACKGROUND

10 Often times it is desirable to search large sets of data, such as
collections of millions of documents, only some of which may pertain to the
information being sought. In such instances it is difficult to either identify a subset
of data to search or to search all data yet return only meaningful results. The
techniques that have been traditionally applied to support searching large sets of
15 data have fallen short of expectations, because they have not been able to achieve
a high degree of accuracy of search results due to inherent limitations.

One common technique, implemented by traditional keyword search
engines, matches words expected to found in a set of documents through pattern
matching techniques. Thus, the more that is known in advance about the
20 documents including their content, format, layout, etc., the better the search terms
that can be provided to elicit a more accurate result. Data is searched and results
are generated based on matching one or more words or terms that are designated
as a query. Results such as documents are returned when they contain a word or
term that matches all or a portion of one or more keywords that were submitted to
25 the search engine as the query. Some keyword search engines additionally
support the use of modifiers, operators, or a control language that specifies how
the keywords should be combined when performing a search. For example, a
query might specify a date filter to be used to filter the returned results. In many
traditional keyword search engines, the results are returned ordered, based on the
30 number of matches found within the data. For example, a keyword search against
Internet websites typically returns a list of sites that contain one or more of the
submitted keywords, with the sites with the most matches appearing at the top of
the list. Accuracy of search results in these systems is thus presumed to be
associated with frequency of occurrence.

One drawback to traditional keyword search engines is that they do not return data that fails to match the submitted keywords, even though the data may be relevant. For example, if a user is searching for information on what products a particular country imports, data that refers to the country as a
5 “customer” instead of using the term “import” would be missed if the submitted query specifies “import” as one of the keywords, but doesn’t specify the term “customer.” For example, a sentence such as “Argentina has been the main customer for Bolivia’s natural gas” would be missed, because no forms of the word “import” are present in the sentence. Ideally, a user would be able to submit a
10 query and receive back a set of results that were accurate based on the meaning of the query – not just on the specific keywords used in submitting in the query.

Natural language parsing provides technology that attempts to understand and identify the syntactical structure of a language. Natural language parsers (“NLPs”) have been used to identify the parts of speech of each term in a
15 submitted sentence to support the use of sentences as natural language queries against data. However, systems that have used NLPs to parse and process queries against data, even when the data is highly structured, suffer from severe performance problems and extensive storage requirements.

Natural language parsing techniques have also been applied to
20 extracting and indexing information from large corpora of documents. By their nature, such systems are incredibly inefficient in that they require excessive storage and intensive computer processing power. The ultimate challenge with such systems has been to find solutions to reduce these inefficiencies in order to create viable consumer products. Several systems have taken an approach to
25 reducing inefficiencies by subsetting the amount of information that is extracted and subsequently retained as structured data (that is only extracting a portion of the available information). For example, NLPs have been used with Information Extraction engines that extract particular information from documents that follow predetermined grammar rules or when a predefined term or rule is recognized,
30 hoping to capture and provide a structured view of potentially relevant information for the kind of searches that are expected on that particular corpus. Such systems typically identify text sentences in a document that follow a particular part-of-speech pattern or other patterns inherent in the document domain, such as “trigger” terms that are expected to appear when particular types of events are
35 present. The trigger terms serve as “triggers” for detecting such events. Other systems may use other formulations for specified patterns to be recognized in the data set, such as predefined sets of events or other types of descriptions of events

or relationships based upon predefined rules, templates, etc. that identify the information to be extracted. However, these techniques may fall short of being able to produce meaningful results when the documents do not follow the specified patterns or when the rules or templates are difficult to generate. The probability of
5 a sentence falling into a class of predefined sentence templates or the probability of a phrase occurring literally is sometimes too low to produce the desired level of recall. Failure to account for semantic and syntactic variations across a data set, especially heterogeneous data sets, has led to inconsistent results in some situations.

10 BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a relationship query and the results returned by an example embodiment of the InFact® 3.1 search engine.

Figure 2 is an example screen display of a custom relationship or event query generator for a website.

15 Figure 3 is an example block diagram of an example Syntactic Query Engine.

Figure 4 is an overview of the steps performed by a Syntactic Query Engine to process data sets and relationship queries.

20 Figure 5 is an example screen display of search results retrieved by relevance and page sorted by action similarity.

Figure 6 is an example screen display of search results retrieved by data and page sorted by date.

Figure 7 is an example screen display illustrating how a user exports search result data.

25 Figure 8 is an example screen display of a report of relationship query result.

Figure 9 is an example screen display of an interface for exporting search result data to a data frame.

30 Figure 10 is an example screen display of a visual interface for specifying attributes of a data frame for export.

Figure 11 is an example screen display of a data frame once exported into another application.

Figure 12 is an example screen display of an interface for a relationship query that specifies a context constraint.

35 Figure 13 is an example screen display of the results of the query specified in Figure 12.

Figure 14 is an example screen display of navigation tips according to a first embodiment.

Figure 15 is an example screen display of the results provided upon selection of a navigation tip according to the first embodiment.

5 Figure 16 is an example screen display of the results provided upon selection of a second navigation tip according to the first embodiment.

Figures 17A and 17B are an example screen display of initial results when a user enters a query based upon a frequently occurring place in a corpus of books.

10 Figures 18A and 18B are an example screen display of results when a user selects a navigation tip.

Figure 19 is an example screen display of results determined when a user selects a deep navigation tip.

15 Figure 20 is an example screen display of the resulting information displayed when a user follows a link to select an individual book.

Figure 21 is an example screen display of example entity instances that correspond to an entity type selected from a book index.

Figure 22 is an example screen display of the resulting information displayed when a user selects a particular entity.

20 Figure 23 is an example screen display of a relationship query executed when a user selects an action pertaining to a selected entity.

Figure 24 is an example screen display of a relationship query executed when a user selects another entity pertaining to a selected entity.

25 Figure 25 is an example flow diagram of the typical steps performed by an SQE tip engine to process tips.

Figure 26 is an example flow diagram of processing performed by a tip searcher component to determine navigation tips for an indicated relationship structure.

30 Figure 27 is an example flow diagram of processing of a relationship query to generate deep tip results.

Figure 28 is a conceptual block diagram of the components of an example embodiment of an enhanced Syntactic Query Engine.

35 Figure 29 is an example block diagram of a general purpose computer system for practicing embodiments of an enhanced Syntactic Query Engine.

DETAILED DESCRIPTION

It is often desirable to search large sets of unstructured data, such as collections of millions of documents, only some of which may pertain to the information being sought. Traditional search engines approach such data mining typically by offering interactive searches that match the data to one or more keywords (terms) using classical pattern matching or string matching techniques. At the other extreme, information extraction engines typically approach the unstructured data mining problem by extracting subsets of the data, based upon formulations of predefined rules, and then converting the extracted data into structured data that can be more easily searched. Typically, the extracted structured data is stored in a relational database management system and accessed by database languages and tools. Other techniques, such as those offered by Insightful Corporation's InFact® products, offer greater accuracy and truer information discovery tools, because they employ generalized syntactic indexing with the ability to interactively search for relationships and events in the data, including latent relationships, across the entire data set and not just upon predetermined extracted data that follows particular syntactic patterns. InFact®'s syntactic indexing and relationship searching uses natural language parsing techniques to grammatically analyze sentences to attempt to understand the meaning of sentences and then applies queries in a manner that takes into account the grammatical information to locate relationships in the data that correspond to the query. Some of these embodiments support a natural language query interface, which parses natural language queries in much the same manner as the underlying data, in addition to a streamlined relationship and event searching interface that focuses on retrieving information associated with particular grammatical roles. Other interfaces for relationship and event searching can be generated using an application programming interface ("API"). Insightful Corporation's syntactic searching techniques are described in detail in U.S. Provisional Application Nos. 60/312,385, 60/620,550, and 60/737,446 and U.S. Application Nos. 10/007,299, 10/371,399, and 11/012,089 (Publication Nos. US 2004/0221235, US 2003/0233224, and US 2005/0267871).

The syntactic indexing and relationship and event searching techniques describe therein extend the use of traditional keyword search engines to relationship and event searching of data sets. In summary, the syntactic and semantic information that is gleaned from an enhanced natural language parsing process is stored in an enhanced document index, for example, a form of a term-clause matrix, that is amenable to processing by the more efficient pattern (string)

matching capabilities of keyword search engines. Accordingly, traditional keyword search engines, including existing or even off-the-shelf search engines, can be utilized to discover information by pattern (or string) matching the terms of a relationship query, which are inherently associated with syntactic and semantic
5 information, against the syntactically and semantically annotated terms of sentence clauses (of documents) stored in the enhanced document index. As another benefit, the additional capabilities of such search engines, such as the availability of Boolean operations, and other filtering tools, are automatically extended to relationship and event searching.

10 Relationship and event searching, also described as "syntactic searching" in U.S. Application Nos. 60/312,385, 10/007,299, 10/371,399, and 60/620,550 supports the ability to search a corpus of documents (or other objects) for places, people, or things as they relate to other places, people, or things, for example, through actions or events. Such relationships can be inferred or derived
15 from the corpus based upon one or more "roles" that each term occupies in a clause, sentence, paragraph, document, or corpus. These roles may comprise grammatical roles, such as "subject," "object," "modifier," or "verb;" or, these roles may comprise other types of syntactic or semantic information such as an entity type of "location," "date," "organization," or "person," etc. The role of a specified
20 term or phrase (e.g., subject, object, verb, place, person, thing, action, or event, etc.) is used as an approximation of the meaning and significance of that term in the context of the sentence (or clause). In this way, a relationship or syntactic search engine attempts to "understand" the sentence when a query is applied to the corpus by determining whether the terms in sentences or clauses of the corpus
25 are associated with the roles specified in the pending query. For example, if a user of the search engine desires to determine all events in which "Hillary Clinton" participated in as a speaker, then the user might specify a relationship query that instructs a search engine to locate all sentences/documents in which "Hillary Clinton" is a source entity and "speak" is an action. In response, the syntactic
30 search engine will determine and return indicators to all sentences/clauses in which "Hillary Clinton" has the role of a subject and with some form of the word "speak" (e.g., speaking, spoke) or a similar word in the role of a verb.

For example, Figure 1 shows a relationship query and the results returned by an example embodiment of the InFact® 3.1 search engine. In the
35 InFact® 3.1 product, a user of the search engine can specify a search, called a "Fact Search" for a known "source" or "target" entity (or both) looking for actions or events that involve that entity. The user can also specify a second entity and look

for actions or events that involve both the first and second entity. The user can specify a particular action or may specify a type of action or any action. An entity specified as a source entity typically refers to the corresponding term's role as a subject (or subject-related modifier) of a clause or sentence, whereas an entity
5 specified as a target typically refers to the corresponding term's role as an object (or object-related modifier) of a clause or sentence. An action or event typically refers to a term's role as a verb, related verb, or verb-related modifier. Moreover, instead of a specific entity, the user can specify an entity type, which refers to a tag such as an item in a classification scheme such as a taxonomy. A user can also
10 specify a known action or action type and look for one or more entities, or entity types that are related through the specified action or action type. Many other types and combinations of relationship searches are possible and supported as described in the above-mentioned co-pending patent applications.

In the example user interface shown in Figure 1, a relationship or
15 event query is specified in query field 101 to. The query is entered as "terrorist > attack > London" according to Insightful Corporation's Query Language ("IQL", also known as "RQL" for Relationship Query Language), as described in detail elsewhere. The query specifies a value "terrorist" for a source entity field 101a, a value "attack" for an action field 101b, and a value "London" for a target entity field
20 101c. The source field 101a and target field 101c indicate whether a specified entity is to be a source of the action or a recipient (target) of the action. The directional arrows in the query specify whether action directionality for events. The particular query displayed instructs the search engine to look for sentence clauses that discuss terrorist attacks in London when the Search button 103 is pressed.
25 The results are returned in summary information field 102, which is shown sorted by similarity to the query. Each matching relationship is shown as a row 107 in the results field 102, optionally with duplicates combined. A researcher or other user can select one of the action links in the actions column 105 to look at the document in the corpus where the relationship or event was found.

30 As an alternative to specifying the relationship query using IQL, the user can use a graphical/form-based interface, termed here a "query generator," for example, by selecting the "Try your own Fact Search" link 110. In response, a form such as that illustrated in Figure 2 is presented. Figure 2 is an example screen display of a custom relationship or event query generator for a website.
35 The user can enter a source entity in field 102, a target (or recipient of the action) in field 104, and an action in field 103. In addition, the user can constrain the

results by specifying particular keywords that need to appear. The results of the generated query are displayed similar to that shown in Figure 1.

More details of the relationship query language and of an example user interface provided by Insightful Corporation's products are described in co-
5 pending patent application 11/012,089.

Embodiments described herein provide enhancements to the methods, systems, and techniques for syntactically indexing and searching data sets to achieve more accurate search results with greater flexibility and efficiency than previously available. Example embodiments provide improvements to an
10 enhanced Syntactic Query Engine ("SQE") that parses, indexes, and stores a data set, as well as performs syntactic searching in response to queries submitted against the data set. In one embodiment, the SQE includes, among other components, a data set repository and an Enhanced Natural Language Parser ("ENLP"). The ENLP parses each object in the data set (typically a document) and
15 transforms it into a canonical form (also termed a normalized form) that can be searched efficiently using techniques of the present invention. To perform this transformation, the ENLP determines the syntactic structure of the data by parsing (or decomposing) each data object into syntactic units, determines the grammatical roles and relationships of the syntactic units, associates recognized
20 entity types if configured to do so, and represents these relationships in a normalized form. The normalized data are then stored and/or indexed as appropriate for efficient searching. Thus, the SQE can provide a unified knowledge representation for both structured and unstructured data.

In example embodiments of the SQE described herein, the
25 normalized data, including the grammatical role and other tag information that can be used to discover or explore relationships, are integrated into enhanced versions of document indexes that are typically used by traditional keyword search engines to index the terms of each document in a corpus. A traditional keyword search engine can then search the enhanced indexing information that is stored in these
30 document indexes for matching relationships in the same manner the search engine searches for keywords. That is, the search engine looks for pattern/string matches to terms associated with the desired tag information as specified (explicitly or implicitly) in a query. A detailed description of an example SQE that provides enhanced indexing information according to these techniques is described
35 in U.S. Patent Application No. 11/012,089. In one such example system, the SQE stores the relationship information that is extracted during the parsing and data object transformation process (the normalized data) in an annotated "term-clause

matrix," which stores the terms of each clause along with "tagged terms," which include the syntactic and semantic information that embodies relationship information. Other example embodiments may provide different levels of organizing the enhanced indexing information, such as an annotated "term-
5 sentence matrix" or an annotated "term-document matrix." Other variations of storage organization are possible, including that each matrix may be comprised of a plurality of other data structures or matrices.

The integration of the enhanced indexing information into traditional keyword search engine type document indexes (for example, an inverted index) is
10 what supports the use of standard keyword search techniques to find a new type of document information – that is, relationship information – easily and quickly. An end user, such as a researcher, can pose simple Boolean style queries to the SQE yielding results that are based upon an approximation of the meaning of the indexed data objects. Because traditional search engines do not pay attention to
15 the actual contents of the indexed information (they just perform string matching or pattern matching operations without regard to the meaning of the content), the SQE can store all kinds of relationship information in the indexed information and use a keyword search engine to quickly retrieve it. In addition, standard document searches can be combined with more specialized relationship searches to
20 generate many types of results.

The SQE processes each query by translating or transforming the query into component keyword searches that can be performed against the indexed data set using, for example, an "off-the-shelf" or existing keyword search engine. These searches are referred to herein for ease of description as keyword
25 searches, keyword-style searches, or pattern matching or string matching searches, to emphasize their ability to match relationship information the same way search terms can be string- or pattern-matched against a data set using a keyword search engine. The SQE then combines the results from each keyword-style search into a cohesive whole that is presented to the user.

30 Figure 3 is an example block diagram of an example Syntactic Query Engine. A document administrator 302 adds and removes data sets (for example, sets of documents), which are indexed and stored within a data set repository 304 of the SQE 301. When used with keyword style searching techniques, the data set repository 304 stores an enhanced document index as described above. In the
35 example shown in Figure 3, a subscriber 303 to a document service submits queries to the SQE 301, perhaps using a visual interface such as a web page. The queries are then processed by the SQE 301 against the data sets indexed in

the data set repository 304. The query results are then returned to the subscriber 303. In this example, the SQE 301 is shown implemented as part of a subscription document service, although one skilled in the art will recognize that the SQE may be made available in many other forms, including as a separate application/tool, integrated into other software or hardware, for example, cell phones, personal digital assistants ("PDA"), or handheld computers, or associated with other types of existing or yet to be defined services. Additionally, although the example embodiment is shown and described as processing data sets and queries that are in the English language, one skilled in the art will recognize that the SQE can be implemented to process data sets and queries in any language, or any combination of languages.

Figure 4 is an overview of the steps performed by a Syntactic Query Engine to process data sets and relationship queries. Steps 401-405 address the indexing (also known as the ingestion) process, and steps 406-409 address the query process. Note that although much of the discussion herein focuses on ingestion of an entire data set prior to searching, the SQE also handles incremental document ingestion. Also, the configuration process that permits an administrator to set up ontologies, dictionaries, sizing preferences for indexes and other configuration and processing parameters is not shown.

Specifically, in step 401, the SQE receives a data set, for example, a set of documents. The documents may be received electronically, scanned in, or communicated by any reasonable means. In step 402, the SQE preprocesses the data set to ensure a consistent data format. In step 403, the SQE parses the data set, identifying entity type tags and the syntax and grammatical roles of terms within the data set as appropriate to the configured parsing level. For the purpose of extending keyword searching to syntactically and semantically annotated data, parsing sufficient to determine at least the subject, object, and verb of each clause is desirable to perform syntactic searches in relationship and event queries. However, as described elsewhere in co-pending patent applications, subsets of the capabilities of the SQE could be provided in trade for shorter corpus ingestion times if full syntactic searching is not desired. In step 404, the SQE transforms the each parsed clause (or sentence) into normalized data by applying various linguistic normalizations and transformations to map complex linguistic constructs into equivalent structures. Linguistic normalizations include lexical normalizations (e.g., synonyms), syntactic normalizations (e.g., verbalization), and semantic normalizations (e.g., reducing different sentence styles to a standard form). These heuristics and rules are applied when ingesting documents and influence how well

the stored sentences eventually will be “understood” by the system. In step 405, the SQE stores the parsed and transformed sentences in a data set repository. As described above, when the SQE is used with a keyword search engine, the normalized data is stored in (used to populate) an enhanced document index such as the term-clause matrix.

After storing the data set, the SQE can process relationship queries against the data set. In step 406, the SQE receives a relationship query, for example, through a user interface such as that shown in Figure 1. Alternatively, one skilled in the art will recognize that the query may be transmitted through a function call, batch process, or translated from some other type of interface. In step 407, if necessary (depending upon the interface) the SQE preprocesses the received relation query and transforms it into the relationship query language understood by the system. For example, if natural language queries are supported, then the natural language query is parsed into syntactic units with grammatical roles, and the relevant entity and action terms are transformed into the query language formulations understood by the SQE. In step 408, the SQE executes the received query against the data set stored in the data set repository. The SQE transforms the query internally into sub-queries as appropriate to the organization of the data in the indexes and executes a traditional keyword search engine (or its own version of keyword style searching) to process the query. In step 409, the SQE returns the results of the relationship query, for example, by displaying them through a user interface such as the summary information 102 shown in Figure 1.

Although the techniques are described primarily with reference to text-based languages and collections of documents, similar techniques may be applied to any collection of terms, phrases, units, images, or other objects that can be represented in syntactical units and that follow a grammar that defines and assigns roles to the syntactical units, even if the data object may not traditionally be thought of in that fashion. Examples include written or spoken languages, for example, English or French, computer programming languages, graphical images, bitmaps, music, video data, and audio data. Sentences that comprise multiple words are only one example of a phrase or collection of terms that can be analyzed, indexed, and searched using the techniques described herein. One can modify the structures and program flow exemplified herein to account for differences in types of data being indexed and retrieved. The concepts and techniques described are applicable to any environment where the keyword style searching is contemplated.

Also, although certain terms are used primarily herein, one skilled in the art will recognize that other terms could be used interchangeably to yield equivalent embodiments and examples. In addition, terms may have alternate spellings which may or may not be explicitly mentioned, and one skilled in the art will recognize that all such variations of terms are intended to be included. For example, the terms "matrix" and "index" are used interchangeably and are not meant to imply a particular storage implementation. Also, a document may be a single term, clause, sentence, or paragraph or a collection of one or more such objects. Also, when referring to various data, aspects, or elements in the alternative, the term "or" is used in its plain English sense, unless otherwise specified, to mean one or more of the listed alternatives.

As additional examples, the term "query" is used herein to include any form of specifying a desired relationship query, including a specialized syntax for entering query information, a menu driven interface, a graphical interface, a natural language query, batch query processing, or any other input (including API function calls) that can be transformed into a Boolean expression of terms and annotated terms. Annotated terms are terms associated with syntactic or semantic tag information, and are equivalently referred to as "tagged terms." Semantic tags include, for example, indicators to a particular node or path in an ontology or other classification hierarchy. "Entity tags" are examples of one type of semantic tag that points, for example, to a type of ENTITY node in an ontology. In addition, although the description is oriented towards parsing and maintaining information at the clause level, it is to be understood that the SQE is able to parse and maintain information in larger units, such as sentences, paragraphs, sections, chapters, documents, etc., and the routines and data structures are modified accordingly. Thus, for ease of description, the techniques are described as they are applied to a term-clause matrix. One can equivalently apply these techniques to a term-sentence matrix or a term-document matrix.

In the following description, numerous specific details are set forth, such as data formats and code sequences, etc., in order to provide a thorough understanding of the described techniques. The embodiments described also can be practiced without some of the specific details described herein, or with other specific details, such as changes with respect to the ordering of the code flow, different code flows, etc. Thus, the scope of the techniques and/or functions described are not limited by the particular order, selection, or decomposition of steps described with reference to any particular routine.

As described in co-pending U.S. Patent Application No. 11/012,089, published as U.S. Patent Publication No. 2005/0267871, many different interfaces are possible for specifying relationship queries, including form based (template-style) interfaces, relationship search strings, visual GUI drag and drop interfaces, a query language called IQL, an application programming interface ("API") for creating application access to relationship query information or for creating customized (also alternative) user interfaces etc. Many of these enhancements are exemplified in Insightful Corporation's InFact® System Release 3.1.

Embodiments of the interfaces include various enhancements to the extended keyword searching methods, systems, and techniques described herein, including, amongst other features, support for a new context operator; improved metadata searching performance, support for ontology (entity tag) searches at the document search level, an improved interface for exporting relationship data, data sorting across an entire result, and automatically generated navigation tips.

15 Context Operator

Embodiments of the enhanced SQE provide a context operator for determining in a relationship search whether the prescribed relation might be satisfied by searching surrounding sentences (clauses, or any other granularity desired, etc.). In some embodiments, the number of surrounding sentences/clauses searched is a predetermined number "n." In other embodiments, it is contemplated that this number can be set dynamically, even from within the query itself. Combinations and permutations (such as permitting entry of any number less than "n") are also possible.

The context operator (also termed a context constraint) can be specified in a query as:

Bush > visit > [Country] AND NOT China

PREP CONTAINS plane

CONTEXT CONTAINS "foreign service" OR diplomat

In this example, a match is found if the terms "foreign service" or "diplomat" are found anywhere within one sentence (before or after) of the sentence matching the rest of the relationship specification. An abbreviated form of the context operator ("~") can also be used:

Bush > visit > [Country] AND NOT China

^ plane

~ "foreign service" OR diplomat

Note that here context is defined as within a sentence. Other alternatives, such as two sentences, clauses, etc. could alternatively be defined. The same techniques can be used to provide an interface at the these levels.

The context operator is particularly beneficial to catch related text
 5 when two separate clauses (sentences) that have no inter-clause connections contain information that is related to what the researcher is searching for. For example, the sentences:

The DC-8 crashed in Japan, after two engines failed. Two hundred people were killed.

10 contain information about a plane crash where the number of people killed appears in a separate sentence. If one is searching to find all documents where a plane crashed injuring or killing some number of people, there are many possible ways to define a query. In one instance, suppose that the SQE defines an entity type such as "plane." Then, if the query is specified as:

15 Query 1: *[plane] > killed > [numeric]*

to avoid getting all sentences where some number of people have been killed regardless of the actor, then neither sentence above would be found. This would also be the case if the query specified "DC-8" instead of an entity type. However, by changing the query to a broader query but constrained by a context operator,
 20 such as:

Query 2: ** > killed > [numeric] ~[plane]*

then, the second sentence above indeed will be found. Other examples abound.

Figures 12 and 13 provide another example of the context operator. Figure 12 is an example screen display of an interface for a relationship query that
 25 specifies a context constraint. In this example, the researcher is interested in finding out information regarding terrorist attacks in London where Osama Bin Laden may have been involved. Accordingly, the query entered into query entry field 1201 specifies "terrorist > attack > London ~bin laden." The results of the query are shown in the result summary area 1202. To view the resulting data in
 30 context, the researcher selects the link 1203. Figure 13 is an example screen display of the results of the query specified in Figure 12. A portion of the corresponding document 1301 shows the matching sentence 1302 that matches the query. Note that the phrase "bin laden" 1303 appears in the context of the sentence following the matching sentence 1302. If instead of using a context
 35 operator the researcher had indicated that "bin laden" need appear somewhere in

the document (for example, using the document constraint operator), many more results may have appeared where Osama Bin Laden was not necessarily discussed in the context of a London attack. Accordingly, by using the context operator, the researcher is able to discover information that may have taken
 5 searching through many results to find.

In order to support context operators, the data structures used to stored the normalized (extended keyword) data have been modified. In summary, the context, stored as a "bag of words," is stored in the term-clause index for fast access to context. In addition, since context information is now searched for
 10 surrounding sentences, the ontology paths present in the surrounding sentences (all entities) are stored for quick comparisons. See the "context" and "context_ont_path" fields in the modified term-clause index described in Table 1 below. The term-clause index and other relevant data structures were described in detail in co-pending U.S. Patent Application No. 11/012,089.

15 Metadata Constraint Performance

Embodiments of an enhanced SQE provide improved techniques for efficiently handling meta data used in relationship searches. Metadata filtering allows one to constrain a search based on document level metadata constraints. Metadata may or may not be available for a particular document corpus. (A user
 20 can determine whether such metadata is available and the various types of information available by selecting a link to the "Corpus" page as available through some embodiments of Insightful Corporation's SQE.) For example, the following (equivalent) queries, search for all documents having at least one sentence indicating Clinton in a "visits" relationship to China, where the author of the
 25 document is "John Smith."

Clinton > visit > China METADATA CONTAINS Author="John Smith"

or

Clinton > visit > China # Author="John Smith"

Such a search might be useful for example, if one was searching for a book by a
 30 known author on a particular subject, but one didn't remember the title.

In the enhanced SQE, the data structures used to stored the normalized (extended keyword) data have been modified to more efficiently store metadata, and specific metadata such as a document date, for fast retrieval. In one example embodiment, the term-document index is modified to contain
 35 additional fields such as "docdate" and "author." In some embodiments, all of the

metadata is available as an inverted index as well. See also, Table 2 described below.

Ontology Specifications in Document level searches

Embodiments of the SQE have been enhanced to support ontology path specifications in document level searches, regardless of whether the search is performed as part of a document (*i.e.*, keyword) search or relationship search. For example, one can search for all documents which contain a specific entity type or ontology path the same way one specifies a keyword search. For example, the general document search query:

10 "DC-8"

will find all documents where a "DC-8" is described somewhere in the document (including the above sentences used in the context operator example). The query can be modified by specifying a single ontology path or entity type such as:

[plane]

15 to achieve slightly broader results. Also, Boolean operators can be used to combine the ontology specifications in the document level search, for example:

"DC-8" AND [country] AND [person] AND NOT [money]

will find all documents that involve a DC-8, a country, and a person as long as the document doesn't contain a reference to money. This would eliminate, for example, results that describe buying a plane in China, but include results where a
20 DC-8 killed someone in China.

One can also use relationship searches with document constraints (also called document specifications) to search for documents with a specific keyword in it. For example, the relationship search exemplified in Query 2 above
25 can be modified to find documents where a "DC-8" is described somewhere in the document (and not just in the surrounding sentences) as:

* > kill > [numeric]

DOCUMENT CONTAINS "DC-8"

or as:

30 * > kill > [numeric]; "DC-8"

Using ontology path (or entity) document level enhancements, the above example relationship query can be modified to find killings of any number of entities happening in any country by:

* > kill > [numeric]; [country]

- 5 Other specifications are of course possible.

The data structures used to stored the normalized (extended keyword) data have been modified to efficiently support document level searches for ontology path information by including such information in specific fields in, for example, the term-document index. In one example embodiment, all ontology
10 paths (and sub-paths) of each available entity are stored in the term-document index. See Tables 1-3 below.

Note as well that the ontologies themselves can be stored in inverted indices for easy searching and access by the SQE. Specifically, ontology entity data can be stored in keyword search engine compatible indices such as those
15 illustrated in Appendix C, which is incorporated herein by reference in its entirety.

Data Sorting Across Entire Result

Another feature of the enhanced SQE is that it can provide search results organized by data across the whole result, and not just on a per page basis. Figure 5 is an example screen display of search results retrieved by
20 relevance and page sorted by action similarity. A summary of the search results is shown in summary result area 503. Because the SQE may return a large amount of results, they are brought in as data "chunks" and each page can be sorted. To sort on a per page basis, the user selects the page sort input field 502. In one embodiment, different page sorting is available depending upon the mode in which
25 the search results are retrieved (globally across the result). In order to change the mode in which the search results are retrieved, the user selects the "retrieve by" field 501. In one embodiment, this field toggles between "retrieve by relevance" and "retrieve by date." Other modes are contemplated and could be similarly incorporated.

30 When the user selects the retrieve by date field 501, the results are displayed organized by date. Figure 6 is an example screen display of search results retrieved by data and page sorted by date. The search results are shown in summary result area 603 arranged by date. The user can change the sorting on a per page basis by selecting a different criterion in page sort input field 601 and

selecting the "sort" button 602. To return to displaying the results organized by relevance, the user selects the "retrieve by relevance" toggle 604.

Exporting Relationship Data

The ability to share relationship data with other applications is useful, especially when there are large amounts of data to be analyzed and statistical tools could be applied to model the data. The enhanced SQE provides an improved interface for exporting relationship query results in a (character delimited) format that can be input to analytic tools such as those provided by S-PLUS or Excel.

Figure 7 is an example screen display illustrating how a user exports search result data. In Figure 7, the SQE provides a view/export control 701 for viewing reports or tables and for exporting the same. The user selects whether to view or export via user interface ("UI") control 702. The control 703 is used to indicate whether a report or a file (table) is desired. Once the user has made these selections, the user presses the "go" button 704 to perform the view/export function.

Figure 8 is an example screen display of a report of relationship query result. Each result 801 and 802 is displayed with relevant information such as the sentence that matched the query and the terms that matched the relationship query components (sources, actions, targets).

Figure 9 is an example screen display of an interface for exporting search result data to a data frame. The user selects "export" via UI control 901, "data frame" via UI control 902, and the "go" button 903. In response, the SQE displays a visual interface for specifying the various columns for the data frame (file).

Figure 10 is an example screen display of a visual interface for specifying attributes of a data frame for export. Using visual interface 1001, the user selects what components of a relationship query, including tags such as grammatical role tags and/or entity tags, metadata, etc. should be assigned to which columns. As shown the user has selected the relationship query's "source" term 1002 as column 1; the query's "action" term 1003 as column 2; the query's "target" term 1004 as column 3; a prepositional phrase term 1005 if present in the matching result as column 4; a publishing date (e.g., metadata) 1006 if present as column 5; and the country of an creator of the work (e.g., metadata) 1007 if present as column 6. New columns can be added by selecting the "Add Column" button 1008. A preview 1009 of the results is also available. Once the user is

satisfied with the data frame format to be exported, the user can select the "Export Data" button 1010.

When the data is exported, it is typically stored as a tab-delimited file. Different delimiters can be incorporated through configuration parameters. Figure 11 is an example screen display of a data frame once exported into another application. For example, the application illustrated in Figure 11 is S-PLUS, and the relationship query result data exported using the SQE has been imported into an S-PLUS object (displayed as 1101) for further analysis and modeling using standard S-PLUS statements to import data into an S-PLUS variable.

10 Automatic Guided Navigation Tips

Embodiments of syntactic query engines have been enhanced to incorporate automated tips for users to increase the ability of a user to specify a desired search. Based upon the inverted indexes used to represent the corpus, the ontologies used in the system and other rules (such as popularity of term choice, etc.) these tips act as "suggestions" to the user to specify with greater particularity a possibly desired search.

Behind the scenes, when a query (*e.g.*, either a document level search such as one or more keywords or a relationship search, including portions of a or an entire relationship query using IQL) is entered, the SQE attempts to determine what relationship searches might be desired based upon a set of rules parsed and evaluated by the SQE. Then, the SQE attempts to run some number of these potential searches in the background (up to all such possibilities). Thereafter, depending upon the particular tip user interface, some indication of the navigation tips and/or search results of the tips are presented to the user. Once displayed, the user can then choose one or more of these preformed searches (by selecting the navigation tips) to quickly see results.

The query may take the form of either one or more keywords, an IQL expression, or components of an IQL expressed query. For example, a user may enter the keywords "Japan China" to try to determine relationship and event information relating to both countries, and the SQE will respond with appropriate navigation tips to assist the user to discover more information. As another example, the user might type in a partial IQL expression such as "China <> * ~ Japan", and the SQE will attempt to recognize the input as particular components of IQL such as a source entity and a context operator expression.

The rules used to determine the navigation tips are specified, typically in a configuration file, by an administrator when the searching system is

configured (or at other times). Each rule specifies some sort of template or specification for determining what relationship query to execute based upon the recognition of particular input. For example, a rule might specify that when a single entity is recognized as input, a rule that treats it as a source entity and looks for related organizations is fired. Such a rule might be expressed, for example, by

single entity $\rightarrow$ "\$IQLSource <> * <> [organization]"

where "\$IQLSource" represents the source entity component of an IQL expression. Examples of different types of rule mapping specifications are illustrated in Appendices A and B, which are herein incorporated in their entirety.

Different heuristics and/or rules can be used to determine which potential search alternatives to suggest, the order they are suggested, etc. For example, the number of results returned, the popularity of a particular search (e.g., how many times it has been executed against that corpus), or other measurements of value or interest can be incorporated. Note that other types of rules and an entire programming or scripting language for running and combining potential queries can be defined for use by the SQE.

In one embodiment, the searches are run to validate whether there are any results for a particular tip before showing them to a user. In another embodiment, the searches are run (typically in the background and even in parallel), the results cached, and indications to some number of them presented to the user. In some embodiments, the cached tip results are first consulted to speed up response time. The cached tip results can also be useful to return tips to a user when an SQE facility is under heavy load and desires to suspend tip rule processing. Also, depending upon the architecture and infrastructure used to evaluate and perform the tip searches (which execute typically in the background), the SQE can perform load balancing, parallel processing, etc.

According to one example embodiment, which can be accessed via url "www.globalsecurity.org," the SQE provides navigation tips to a user based upon whether it can recognize the specification of a (single) entity that is part of a configured ontology. More specifically, the SQE looks at the user specified entity and determines whether there are one or more ontology paths that include the user-specified term as a "leaf" node (an entity). If not, no tips are suggested. If so, then the SQE runs one or more appropriate relationship searches that corresponds to each possible ontology path, in the background using rules that define which searches to run in order to "validate" the searches against the particular corpus. In some instances, only searches that result in matches are then presented to the user.

In some implementations, a configuration file containing rules is supplied (for example by a system administrator) and a mapping (or other means for specifying a set of rules) is stored between entities to be recognized and potential searches to execute. The configuration file may be parsed and the mappings stored when the system comes up or at other times. Appendix A contains one example of a set of ontology path specifications and mappings to corresponding queries as specified for such a configuration file. Also, depending upon the implementation, the SQE may implement different orders or precedence for applying recognized rules. In one case, the rules are applied in the order that they are encountered in the configuration file, so if an entered query matches multiple rules in the file, the first encountered set of rules will apply. In another case, all of the matched rules are applied and all of the potential searches performed in the order specified. Other orders of precedence can be similarly incorporated.

Figure 14 is an example screen display of navigation tips according to a first embodiment. In Figure 14, the user enters a term (e.g., one or more keywords) or entity in the query field 1401 (for example "Bush"), and the SQE presents possible tips to the user in tip link area 1402. In the illustrated example, the SQE has determined (after running appropriate queries in the background and according to a rules file) that the tips shown as tip links 1403 are viable relationship searches for that particular corpus. Internally, the SQE may apply a series of rules and/or heuristics to determine whether a particular relationship search is viable; such as, a minimum/maximum number of results returned, the top "n" in a ranked ordering, etc. The tip link area 1402 informs the user a little about the search, using the general structure "View facts involving <entity name> and <tip1> <tip2> ... <tip n>," where the <entity name> and <tip> are appropriately substituted.

To use a tip, the user selects one of the tips from the tip link area 1402. Figure 15 is an example screen display of the results provided upon selection of a navigation tip according to the first embodiment. In Figure 15, the user has (previously) indicated the entity "Bush" (see entity name 1501) and has selected tip link 1502 "US Officials," because the user is interested in seeing recent news on the resignation of President Bush's Secretary of Defense, Donald Rumsfeld. The results of the tip are shown in summary result area 1503.

Alternatively, suppose that the user wanted to see who President Bush chose as a replacement Secretary of Defense – the user knew the person came from another government agency, but couldn't remember which agency. Figure 16 is an example screen display of the results provided upon selection of a

second navigation tip according to the first embodiment. In Figure 16, the user has selected the "Organizations" tip link 1602 and can view a summary of the results in summary area 1604. Note that from the summaries alone shown in Figures 15 and 16, the user can quickly get the desired information. Then, if more detail is
5 desired, the user can follow the action link (e.g., "call" 1605) as described elsewhere to view the underlying document.

Note as well that different user interfaces can be used to display the tips. In one such interface, a menu of possible search specifications is provided. In another interface, links are presented that can be selected to show "fast search"
10 results.

According to a second example embodiment, which can be accessed via url "books.infact.com," the SQE provides navigation tips to a user based upon whether it can recognize one or more keywords (as in a document level search) or the specification of (any part of) a relationship query using IQL. More specifically,
15 the SQE looks at the user specified entry and determines whether there are rules that map additional relationship queries to the keywords and/or recognized (portion of a) relationship query. In the case of keyword input, the SQE first parses the input (using the extended natural language parser of the SQE) to transform the input into a relationship query structure (an enhanced or normalized data structure
20 of the SQE) that can be compared with the rules. Appendix B contains one example of a configuration file having a set of mappings from a relationship query structure (which may contain a single entity or action) to corresponding additional queries. Note that the configuration file of Appendix B specifies rules using XML tag definitions; however, as indicated above, other languages, specifications, and
25 mappings could be used. Based upon evaluating the corresponding rules, the SQE then display the potential other relationship queries that might be of interest to the user. Figures 17-24 are example screen displays from this example navigation tip interface.

Figures 17A and 17B are an example screen display of initial results
30 when a user enters a query based upon a frequently occurring place in a corpus of books. In this example, the user has entered "United Kingdom" (which is displayed in query field 1710) and the SQE lists all of the books relating in some manner with the United Kingdom ("UK") in a query results summary area 1701. Entries for each matching book 1702-1705 are shown, along with links based upon the matching
35 location in the book (where "United Kingdom" or an equivalent synonym, abbreviation, etc.) occurs. The user can follow such a link to see the context in the book where "United Kingdom" (or its equivalent term) appears. The SQE also

presents navigation tips to the user in tips area 1730, which are the links to additional relationship queries that the user can run to obtain more information. In this particular example, the tips are organized in groups (entity types or action types) such as the "organizations" group 1731. Within the organization group 5 1731, a general relationship query 1740 corresponding to [organizations] as a type of entity (e.g., any organization) is shown. In addition, "deep" tips 1741 are provided where particular instances of the entity type relationship search have yielded search results. In some implementations these search results may be cached for quick access. Another general relationship query 1742 (see Figure 10 17B) shows a navigation tip corresponding to [religion] (e.g., any religion) as type of entity. Relationship queries that correspond to specific instances of the religions found in this corpus are displayed as deep tips 1743 as long as they meet whatever presentation rules and heuristics are defined by the SQE tip subsystem (e.g., all queries that yield more than 1 result).

15 Note as well that, in some embodiments, synonyms for tips are collapsed and other normalizations of the tips are performed. For example, if during the tip construction process the SQE determines that relationship queries yield valid results against the corpus for "Bill Clinton," "William Clinton," and "President Clinton," then the tip subsystem may display a single deep tip "Clinton." 20 Similarly, synonyms such as "UK," "Great Britain," and "United Kingdom" will be coalesced into one tip. Other normalizations, such as upper and lower case transformations may also be performed.

Figures 18A and 18B are an example screen display of results when a user selects a navigation tip. Figure 18A shows the results in result summary 25 area 1801 after the user has selected the "any religion" tip 1742 in Figure 17B. Query information area 1810 shows that the relationship query that was executed resulted in books that describe ways in which the United Kingdom relates to religion. Each relationship query that was evaluated by the SQE in response to selection of the tip 1742, for example queries 1820, 1821, and 1822, is illustrated 30 as part of the corresponding matching book entry. Each book entry shows the book's title 1830, the matching sentence in the book, the portions of the sentence that match the relationship query in bold, and the source entity 1820a, action 1820b, and target entity 1820c that were evaluated in the corresponding relationship query 1820.

35 Figure 19 is an example screen display of results determined when a user selects a deep navigation tip. In this case, the user selects the deep tip "France" 1901 to show the relationships between the United Kingdom 1902 and

France 1903 as described by the books in the corpus. By observing the relationship queries (e.g., queries 1904, 1905, and 1906), it is apparent that different relationship queries (*i.e.*, different IQL statements) have been executed that use (recognize) the input (including the user input and the selected entity/action reflected by the tip) to correspond to different aspects of a relationship search. For example, relationship query 1904 shows that "Great Britain" (as a synonym of the "United Kingdom" as an entity) is used in that query as a source entity of the relationship query 1904 and "France" is used as a target entity of the relationship query 1904. Relationship query 1905, however, shows both entities "Great Britain" and "France" as being part of only the target entity. Also different, relationship query 1906 shows "Great Britain" as being used as both a source and target entity. (The term "France" is likely specified as a document constraint or as a context operator, depending upon the specific corresponding rule.) Thus, the relationship searches that are suggested by the tips tend to give different suggestions than might otherwise be readily apparent.

In the books corpus example, a user can also peruse individual books and quickly get more information using relationship searching and navigation tips. Figures 20-24 describe some of the aspects of navigation tips used to discover information about a single book.

Figure 20 is an example screen display of the resulting information displayed when a user follows a link to select an individual book. In this case, the book selected is book 1704 in Figure 17. Information about the entities present in the book 2002 is shown as part of book index area 2003. Results regarding matching text (to the prior relationship query "United Kingdom") are shown in result summary area 2010. Additional navigation tips – this time regarding events (actions) that may be discussed in the books are described in tip area 2004.

When the user selects an entity type in the book index area 2003, the various instances of entities that appear in the book are displayed. Figure 21 is an example screen display of example entity instances that correspond to an entity type selected from a book index. In Figure 21, the user has selected the "character" entity type 2105, and the particular characters that appear in the book are displayed as characters 2106 in result summary area 2101. To select a particular character, the user selects a link to the entity (for example, "Steele" 2107).

Figure 22 is an example screen display of the resulting information displayed when a user selects a particular entity. In this example, two passages 2207 and 2208 describe the character, "Steele" in the selected book. The user

can further explore what Steele does in the book, or events that relate to that character, by selecting one of the action tips 2210.

Figure 23 is an example screen display of a relationship query executed when a user selects an action pertaining to a selected entity. In Figure 5 23, the user has selected the navigation tip for the action "recommend" 2310, and the results 2311 and 2312 of executing a corresponding relationship query are shown. Note as well that a history of the executed relationship queries is available as history control 2320, allowing a user to become more familiar with relationship queries and to re-execute them quickly.

10 Figure 24 is an example screen display of a relationship query executed when a user selects another entity pertaining to a selected entity. In Figure 24, the user has selected the navigation tip for a second entity, in this case an entity type "any character" 2413, to see the relationships between the poet Steele and other people in the book. The results 2415 and 2416 of executing a 15 relationship query that corresponds to the selected tip 2413 are shown.

There are any number of techniques an SQE can use to provide the interfaces and navigation tips described in Figures 14-24. Figures 25 and 26 illustrate typical operations of an example embodiment of a tip engine/subsystem provided by enhanced SQE. Note that, as there are any number of known ways to 20 architecturally configure and implement the components of such as system using for example distributed computing techniques, parallelism, etc., including for example systems that are similar to those described in co-pending U.S. Patent Application 11/012,089.

Figure 25 is an example flow diagram of the typical steps performed 25 by an SQE tip engine to process tips. In step 2501, the tip engine parses the tip rules and stores the resultant mappings (from entity → rules in the first example embodiment or from relationship query → rules in the second embodiment). For example, when the SQE system initializes (or at other times), it reads the navigation tip rule configuration file and instantiates appropriate rule objects that 30 can be invoked when a rule is found to match a specified query. In step 2502, the tip engine receives indication that a user has specified a query – however simple or complex. In step 2503, the tip engine determines a tip searcher component to process the indicated query, and returns to step 2502 to wait for the next input query.

35 Depending upon the architectural implementation, including the various parallel processing and load balancing techniques and components available, one or more components may actually be invoked to perform tip

processing. In one implementation (not shown) a tip searcher manager is invoked by the tip engine to properly load balance and distribute requests to some number of tip searcher components. In addition, as will be described below, a tip searcher may further distribute execution of various relationship queries so that they can be performed in parallel. Other architectures are also possible.

Figure 26 is an example flow diagram of processing performed by a tip searcher component to determine navigation tips for an indicated relationship structure. In step 2601, the tip searcher receives an indication of an input query as a tip request. In step 2602, the tip searcher determines whether the input query is keywords or already is in the form of a relationship query structure (the input processing performed this). If the input is keywords, then the tip searcher progresses to step 2603 to first parse the input using the Enhanced Natural Language Parser (the "ENLP") and then continues in step 2604. In step 2604, the tip searcher "parses" the relationship query structure into its components (e.g., source entity, target entity, and action) and determines from the stored rules which rules apply. As stated above, in some implementations, the rules are applied from top to bottom of a corresponding rule configuration file. In steps 2605-2608 the tip searcher performs a loop to execute the one (or more) rules that apply to the indicated relationship query structure components. Specifically, in step 2605, the tip searcher determines whether there are more rules to execute and, if so, continues in step 2606, otherwise it is done processing. In step 2606, the tip searcher determines from the rule a relationship query (typically a general one, for example, "\$IQLSource <> * <> [organization]" to execute. Then, in step 2607, the tip searcher invokes a subcomponent/routine to execute the rule (perform the indicated relationship query) to get results and determine any deep tips. Of note, this processing could also be performed by the tip searcher itself, and, although Figure 26 shows rule execution as a "loop," in some architectures each rule is processed in parallel. Other organizations and other architectures are applicable.

As an example, the rule file may specify a more general relationship query that involves an entity type or an action type. Once the rule is executed and the general relationship query evaluated, the tip searcher (or its components) can evaluate (weigh, score, filter, etc.) the particular entities/actions discovered as results from evaluating the general relationship query and then execute specific relationship queries for those entities/actions. The specific relationship queries that correspond to entities or actions (as opposed to types of entities or actions) can be used to provide the deep tips described in the screen displays above.

Assuming, for example, that the deep tips are determined asynchronously, at some point the results from executing the general relationship query are made available and the tip searcher is notified. In step 2608, the tip searcher determines from these results whether the general tip has generated
5 results that satisfy the tip system rules and heuristics, and if so, presents them to the user. For example, if a (general) tip generates less than some threshold number of results, the tip system may not present a general tip. In other embodiments, the tip searcher always presents the general tip as it corresponds to an available rule. The tip searcher then returns to step 2605 to process the next
10 rule, if any. If not, the tip searcher process is done (until invoked again).

Figure 27 is an example flow diagram of processing of a relationship query to generate deep tip results. and presents the corresponding relationship queries as "deep" tips. In summary, the routine executes a general relationship query (using cached results if any are available) and determines from the specific
15 entities and/or specific verbs found in the results whether a corresponding more specific relationship query should be presented as a "deep" tip. Different evaluation processes can be incorporated in making this determination, including ranking the number of times an entity/verb appears and only presenting the top "n" in the rank, determining whether the entity/verb appears more than a threshold
20 number of times, etc. Note that this part of all of this processing could also be performed by the tip searcher. Also, although shown as executing in a loop, the process of processing each entity/verb to obtain a specific relationship search result can be performed in parallel.

Specifically, in step 2701, the routine receives the general
25 relationship query to execute and evaluate. (Note that, in other embodiments, the routine might just receive the rule and be responsible for determining the corresponding relationship query.) In step 2702, the routine determines whether a cached result of the general relationship query is already available and, if so, retrieves it (step 2703); otherwise in step 2603 executes the general relationship
30 query that corresponds to the rule to determine a result. In step 2604, the routine returns the results to the tip searcher (or stores them and notifies the tip searcher) for presentation of a "general" tip, and in steps 2705-2707 processes and evaluates the specific entities and/or actions found in the results.

More specifically, in step 2705 the routine determines whether there
35 is another entity/action to process and, if so, continues in step 2706, other continues in step 2708. In step 2706, the routine processes the next entity/action. In step 2707, the routine evaluates the search results that correspond to this

entity/action to determine whether they meet the criteria associated with presenting a deep tip. This evaluation, as mentioned, may rank all of the results to determine the most frequently appearing entities/actions and/or may determine whether the number of times an entity/verb appears reaches or surpasses some threshold.

5 Other evaluation criteria may of course be incorporated. In any case, in step 2707 once the routine determines whether an entity/action qualifies for a deep tip, it determines a corresponding relationship query using that entity/action and stores the corresponding RQ result. In step 2708 the routine returns relationship queries that correspond to the stored (or "n" number of them) results and causes them to

10 be presented as deep tips, and then finishes processing. Note that the number of results stored/returned may be a predetermined number, a settable number, a default, etc. In addition, a preference variable may be available to change the number while the SQE is running.

Enhanced SQE Architecture

15 Figure 28 is a conceptual block diagram of the components of an example embodiment of an enhanced Syntactic Query Engine. The enhanced Syntactic Query Engine 2801 provides the enhancements and improvements described above including the context operator, metadata searching, ontology searching at the document search level, exporting relationship data, data sorting,

20 and navigation tips. A Syntactic Query Engine 2801 comprises a Relationship Query Processor 2810, a Data Set Preprocessor 2803, a Data Set Indexer 2807, an Enhanced Natural Language Parser ("ENLP") 2804, a data set repository 2808, and, in some embodiments, a user interface (or an Applications Programming Interface "API") 2813. The Data Set Preprocessor 2803 converts received data

25 sets 2802 to a format that the Enhanced Natural Language Parser 2804 recognizes. The Enhanced Natural Language Parser ("ENLP") 2804, parses the preprocessed sentences, identifying the syntax and grammatical role of each meaningful term in the sentence and the ways in which the terms are related to one another and/or identifies designated entity and other ontology tag types and

30 their associated values, and transforms the sentences into a canonical form — a normalized data representation. The Data Set Indexer 2807 indexes the normalized data into the enhanced document indexes and stores them in the data set repository 2808. The Relationship Query Processor 2810 receives relationship queries and transforms them into a format that the Keyword Search Engine 2811

35 recognizes and can execute. (Recall that the Keyword Search Engine 2811 may be an external or 3rd party keyword search engine that the SQE calls to execute

queries.) The Keyword Search Engine 2811 generates and executes keyword searches (as Boolean expressions of keywords) against the data set that is indexed and stored in the data set repository 2808. The Keyword Search Engine 2811 returns the search results through the user interface/API 2813 to the requester as Query Results 2812.

Details of the operation of an SQE are provided in co-pending U.S. Patent Application No. 11/012,089. In summary, the enhanced SQE 2801 receives as input a data set 2802 to be indexed and stored. The Data Set Preprocessor 2803 prepares the data set for parsing by assigning a Document ID to each document that is part of the received data set (and sentence and clause IDs as appropriate), performing OCR processing on any non-textual entities that are part of the received data set, and formatting each sentence according to the Enhanced Natural Language Parser format requirements. The Enhanced Natural Language Parser ("ENLP") 2804 parses the data set, identifying for each sentence, a set of terms, each term's tags, including potentially part of speech and associated grammatical role tags and any associated entity tags or ontology path information, and transforms this data into normalized data. The Data Set Indexer 2807 indexes and stores the normalized data output from the ENLP in the data set repository 2808. The data set repository 2808 represents whatever type of storage along with the techniques used to store the enhanced document indexes. For example, the indexes may be stored as sparse matrix data structures, flat files, etc. and reflect whatever format corresponds to the input format expected by the keyword search engine. After a data set (or a portion of a data set) is indexed, a Relationship Query 2809 may be submitted to the enhanced SQE 2801 for processing. The Relationship Query Processor 2810 prepares the query for parsing, for example by splitting the Relationship Query 2809 into sub-queries that are executable directly by the Keyword Search Engine 2811. As explained elsewhere, a Relationship Query 2809 is typically comprised of a syntactic search along with optional constraint expressions. Also, different system configuration parameters can be defined that influence and instruct the SQE to search using particular rules, for example, to include synonyms, related verbs, etc. Thus, the Relationship Query Processor 2810 is responsible for augmenting the specified Relationship Query 2809 in accordance with the current SQE configured parameters. To do so, the Relationship Query Processor 2810 may access the ontology information which may be stored in Data Set Repository 2808 or some other data repository. The Relationship Query Processor 2810 splits up the query into a set of Boolean expression searches that are executed by the Keyword

Search engine 2811 and causes the searches to be executed. The Relationship Query Processor 2810 then receives the result of each search from the Keyword Search Engine 2811 and combines them as indicated in the original Relationship Query 2809 (for example, using Boolean operators). Note that the Relationship Query Processor 2810 may be comprised of multiple subcomponents that each execute a portion of the work required to preprocess and execute a relationship query and combine the results for presentation. The results (in portions or as required) are sent to the User Interface/API component 2813 to produce the overall Query Result 2812. The User Interface Component 2813 may interface to a user in a manner similar to that shown in the display screens of Figures 5-24.

The functions of data set processing (data object ingestion) and relationship query processing can be practiced in any number of centralized and/or distributed configurations of client – server systems. Parallel processing techniques can be applied in performing indexing and query processing to substantial increase throughput and responsiveness. Representative configurations and architectures are described in detail in co-pending U.S. Patent Application No. 11/012,089; however, a variety of other configurations could equivalently perform the functions and capabilities identified herein.

Figure 29 is an example block diagram of a general purpose computer system for practicing embodiments of an enhanced Syntactic Query Engine. The computer system 2901 contains one or more central processing units (CPUs) 2902, Input/Output devices 2903, a display device 2904, and a computer memory (memory) 2905. The enhanced Syntactic Query Engine 2920, including the Query Processor 2906, Keyword Search Engine 2907, Data Set Preprocessor 2908, Data Set Indexer 2911, Enhanced Natural Language Parser 2912, and data set repository 2915, preferably resides in memory 2905, with the operating system 2909 and other programs 2910 and executes on the one or more CPUs 2902. Note that the SQE may be implemented using various configurations. For example, the data set repository may be implemented as one or more data repositories stored on one or more local or remote data storage devices. Furthermore, the various components comprising the SQE may be distributed across one or more computer systems including handheld devices, for example, cell phones or PDAs. Additionally, the components of the SQE may be combined differently in one or more different modules. The SQE may also be implemented across a network, for example, the Internet or may be embedded in another device.

As mentioned, the data structures used to store relationship data have been modified to support the enhancements described herein. Specifically, the data set indexer 2807 in Figure 28 stores the normalized data generated from the data set using one or more data structures to provide the abstraction of a term-
 5 clause matrix, a term-sentence matrix, or a term-document matrix. Any data structure that can be understood by the target keyword search engine being used is operable with the techniques described here. In some embodiments, separate indexes exist for each enhanced document matrix (term-clause, term-sentence, term-document).

10 Table 1 below conceptually illustrates the modifications made to the information that is maintained in an example term-clause index to support the enhanced SQE.

Field Name	Type	Description
documentId	Indexed, stored	document id
sentenceId	Indexed, stored	sentence id
clauseId	Indexed, stored	clause id
subject	tokenized, indexed	contains subjects(s), subject modifiers and entity type(s) for subjects and modifiers. The modifiers should be separated into prefix and suffix. If subject has entity type we also store t_entity (just once). If any modifier has entity type we also store tm_entity (just once). We also store noun phrases recognized by NL parser. These noun phrases are stored with spaces replaced by '\.' The subject field order is: prefix_subject_mod subject suffix_subject_mod Entity_types NLP_noun_phrases.
object	tokenized, indexed	contains objects(s), object modifiers and entity type(s) for objects and modifiers. The modifiers should be separated into prefix and suffix. If object has entity type we also store t_entity (just once). If any modifier has entity type we also store tm_entity (just once). We also store noun phrases recognized by NL parser. We also store noun phrases recognized by NL parser. These noun phrases are stored with spaces replaced by '\.' The object field order is: prefix_object_mod object suffix_object_mod Entity_types NLP_noun_phrases.

pcomp	tokenized, indexed	contains pcomp(s), preposition(s), pcomp modifiers and entity type(s) for pcomp, modifiers. The modifiers should be separated into prefix and suffix. If pcomp has entity type we also store t_entity (just once). If any modifier has entity type we also store tm_entity (just once). We also store noun phrases recognized by NL parser. These noun phrases are stored with spaces replaced by '\.' The pcomp field order is: preposition pcomp modifiers, pcomp Entity_types NLP_noun_phrases
verb	tokenized, indexed	contains verbs(s), verb modifiers and entity type(s) for verbs and modifiers. We also store noun phrases recognized by NL parser. These noun phrases are stored with spaces replaced by '\.' The verb field order is: prefix_verb_mod verb suffix_verb_mod Entity_types NLP_noun_phrases.
parent_id	indexed, stored	clause id of a parent clause
clause_rel_sent_class	tokenized, indexed	Contain inter-clause relationships: • conditional_c • causal_c • prepositional_c • temporal_c and Sentence Attributes: • question_s • definition_s • temporal_s • numerical_s.
docdate	Indexed, stored	relationship date, used for ranking/sorting of results
score	Indexed, stored	relationship score, used for ranking/sorting of results
context	tokenized and indexed	Text of the sentence that clause belongs to and neighboring sentences used in context searches.
context_ont_path	tokenized and indexed	Ontology paths for entities in the sentence that clause belongs to and neighboring sentences used in context searches
relationship	stored	Encoded clause for display

Table 1

Table 2 below conceptually illustrates the modifications made to the information that is maintained in an example sentence index to support the enhanced SQE.

Field Name	Type	Description
documentId	indexed	Document id
sentid	indexed	sentence id
sentenceid		Document id, sentence id concatenated with ' '
sent_text	Stored, not indexed	String content of the sentence

5

Table 2

Table 3 below conceptually illustrates the modifications made to the information that is maintained in an example document index to support the enhanced SQE.

10

Field Name	Type	Description
doc_id	Indexed, stored	Document id
dhs_doc_id	Indexed, stored	DHS_doc_id (URL in our case)
title	Indexed, stored	Document title
boost	tokenized and indexed	Important terms used in the document keyword searches
creationDate	Indexed, stored	Document creation date; format: yyyy.MM.dd-HH:mm:ss
document_type	Stored	Format of the document
docdate	Indexed, stored	document date for sorting format: yyyyMMddHH
metatag	Tokenized, Indexed, stored	All metatags stored as separated MetatagName#MetatagValue with spaces in names or values replaced by "/." The type of the metatag is appended to each metatag
ontology_path	Tokenized, Indexed, Not Stored, not stemmed	All detected Ontology paths (Entity types) prefixed by "t_".
content	Tokenized, Indexed, Not Stored, stemmed	String content of the document

Table 3

From the foregoing it will be appreciated that, although specific embodiments have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the present disclosure. In addition, those skilled in the art will understand how to make changes and modifications to the methods and systems described to meet their specific requirements or conditions. The methods and systems discussed herein are applicable to differing protocols, communication media (optical, wireless, cable, etc.) and devices (such as wireless handsets, electronic organizers, personal digital assistants, portable email machines, game machines, pagers, navigation devices such as GPS receivers, etc.). For example, the methods and systems described herein can be applied to any type of search tool or indexing of a data set, and not just the enhanced SQE described. In addition, the techniques described may be applied to other types of methods and systems where large data sets must be efficiently reviewed. For example, these techniques may be applied to Internet search tools implemented on a PDA, web-enabled cellular phones, or embedded in other devices. Furthermore, the data sets may comprise data in any language or in any combination of languages. In addition, the user interface and API components described may be implemented to effectively support wireless and handheld devices, for example, PDAs, and other similar devices, with limited screen real estate. These and other changes may be made in light of the above-detailed description.



Example Entity Guided Navigation Tips

Appendix A

```

#entity|name|query
nuclear weapons|Proliferation|%<>combat OR proliferate OR end OR spread
OR function OR fund
nuclear weapons|Terrorism|%<>*<>*[terrorist]
nuclear weapons|Its Usage/Operation|%<>"use" OR "operate"<>*
nuclear weapons|People|%<>*<>[person/name]
nuclear weapons|Locations|%<>*<>[location]
nuclear weapons|Money|%<>*<>[money]
nuclear arms|Proliferation|%<>combat OR proliferate OR end OR spread OR
function OR fund
nuclear arms|Terrorism|%<>*<>*[terrorist]
nuclear arms|Its Usage/Operation|%<>"use" OR "operate"<>*
nuclear arms|People|%<>*<>[person/name]
nuclear arms|Locations|%<>*<>[location]
nuclear arms|Money|%<>*<>[money]
biological weapons|Proliferation|%<>combat OR proliferate OR end OR
spread OR function OR fund
biological weapons|Terrorism|%<>*<>*[terrorist]
biological weapons|Its Usage/Operation|%<>"use" OR "operate"<>*
biological weapons|People|%<>*<>[person/name]
biological weapons|Locations|%<>*<>[location]
biological weapons|Money|%<>*<>[money]
chemical weapons|Proliferation|%<>combat OR proliferate OR end OR spread
OR function OR fund
chemical weapons|Terrorism|%<>*<>*[terrorist]
chemical weapons|Its Usage/Operation|%<>"use" OR "operate"<>*
chemical weapons|People|%<>*<>[person/name]
chemical weapons|Locations|%<>*<>[location]
chemical weapons|Money|%<>*<>[money]
WMD|Proliferation|nuclear weapons OR biological weapons OR chemical
weapons<>combat OR proliferate OR end OR spread OR function OR fund
WMD|Terrorism|nuclear weapons OR biological weapons OR chemical
weapons<>*<>*[terrorist]
WMD|Its Usage/Operation|nuclear weapons OR biological weapons OR chemical
weapons<>"use" OR "operate"<>*
WMD|People|nuclear weapons OR biological weapons OR chemical
weapons<>*<>[person/name]
WMD|Locations|nuclear weapons OR biological weapons OR chemical
weapons<>*<>[location]
WMD|Money|nuclear weapons OR biological weapons OR chemical
weapons<>*<>[money]
iraq|Kidnappings|[name]<>kidnap<>iraq
iraq|Suicide Bombings|suicide bomb*<>*<>iraq
iraq|Deployments|[organization] OR [person] >"deploy" OR "position" OR
"roll">iraq
iraq|The Insurgency|insurgency AND iraq<>*
iraq|Elections|election AND iraq<>*
iraq|Reconstruction|reconstruction AND iraq<>*
afghanistan|Kidnappings|[name]<>kidnap<>afghanistan
afghanistan|Suicide Bombings|suicide bomb*<>*<>afghanistan
afghanistan|Deployments|[organization] OR [person] >"deploy" OR
"position" OR "roll">afghanistan
afghanistan|The Insurgency|(insurgency or taliban) AND afghanistan<>*
afghanistan|Elections|election AND afghanistan<>*
afghanistan|Reconstruction|reconstruction AND afghanistan<>*
north korea|Nuclear Weapons|nuclear weapons<>*<>north korea
north korea|Negotiations|north korea<>negotiate<>*
north korea|People|%<>*<>[person/name]

```

north korea|Countries|<>*<>[country]
 north korea|Organizations|<>*<>[organization/name]
 north korea|Money|<>*<>[money]
 iran|Nuclear Weapons|nuclear weapons<>*<>iran
 iran|Negotiations|nuclear OR uranium OR weapon OR
 enrichment<>negotiate<>iran
 iran|People|<>*<>[person/name]
 iran|Countries|<>*<>[country]
 iran|Organizations|<>*<>[organization/name]
 iran|Money|<>*<>[money]
 blair|Places Visited|>visit>[location]
 blair|Offense|>hit OR attack>*<>
 blair|Defense|>defend OR deny>*<>
 oil|Prices|oil OR crude<>*<>[money]
 crude|Prices|oil OR crude<>*<>[money]
 crude oil|Prices|oil OR crude<>*<>[money]
 [terrorist]|Casualties|>kill OR die OR martyr>[numeric]~% AND [location]
 [terrorist]|Attacks Performed|>attack OR fight OR kill OR martyr OR
 assassination OR launch OR claim OR shoot OR explode OR plot OR plan
 [terrorist]|Attacks Received|>attack OR fight OR kill OR martyr OR
 launch OR assassinate>*<>
 [terrorist]|People|<>*<>[person/name] AND NOT %
 [terrorist]|Places|<>*<>[location]
 [terrorist]|Organizations|<>*<>[organization/name] AND NOT %
 [terrorist]|Money|<>*<>[money]
 [weapon]|Combat|<>attack OR kill OR injure OR harm OR combat
 [weapon]|Its Usage/Operation|<>use OR function OR fund OR operate<>*<>
 [weapon]|Locations|<>*<>[location]
 [weapon]|Military Organizations|<>*<>[military]
 [weapon]|Money|<>*<>[money]
 [aircraft]|Combat|<>attack OR kill OR injure OR harm OR combat
 [aircraft]|Aircraft|<>*<>[aircraft] AND NOT %
 [aircraft]|Locations|<>*<>[location]
 [aircraft]|Military Organizations|<>*<>[military]
 [aircraft]|Money|<>*<>[money]
 [us-official]|Attacks|>hit OR attack>*<>
 [us-official]|Defense|>defend OR deny>*<>
 [us-official]|US Officials|<>*<>[us-official] AND NOT %
 [us-official]|Countries|<>*<>[country]
 [us-official]|Organizations|<>*<>[organization/name]
 [us-official]|Money|<>*<>[money]
 [person/name]|People|<>*<>[person/name] AND NOT %
 [person/name]|Places|<>*<>[location]
 [person/name]|Organizations|<>*<>[organization/name]
 [person/name]|Money|<>*<>[money]
 [city]|Security|<>*<>security OR terrorism
 [city]|People|<>*<>[person/name]
 [city]|Cities|<>*<>[city] AND NOT %
 [city]|Organizations|<>*<>[organization/name]
 [city]|Money|<>*<>[money]
 [country]|Hostilities|<>bomb OR attack OR fight OR kill AND NOT end<>*<>
 [country]|Security|<>*<>security OR terrorism
 [country]|Weapons|<>*<>[weapon]
 [country]|People|<>*<>[person/name]
 [country]|Countries|<>*<>[country] AND NOT %
 [country]|Organizations|<>*<>[organization/name]
 [country]|Money|<>*<>[money]
 [location]|Security|<>*<>security OR terrorism
 [location]|People|<>*<>[person/name]
 [location]|Locations|<>*<>[location] AND NOT %
 [location]|Organizations|<>*<>[organization/name]

[location] | Money | %<>*<> [money]
 [government] | Security | %<>*<> security OR terrorism
 [government] | People | %<>*<> [person/name]
 [government] | Places | %<>*<> [location]
 [government] | Organizations | %<>*<> [organization/name] AND NOT %
 [government] | Money | %<>*<> [money]
 [trade] | Agreements | %<>*<> agreement OR accord OR consensus
 [trade] | Discussions | %<>discuss<>*
 [trade] | People | %<>*<> [person/name]
 [trade] | Places | %<>*<> [location]
 [trade] | Organizations | %<>*<> [organization/name] AND NOT %
 [trade] | Money | %>*> [money]
 [military] | Deployments | %<>deploy OR depart OR leave
 [military] | Security | %<>*<> security OR terrorism
 [military] | People | %<>*<> [person/name]
 [military] | Places | %<>*<> [location]
 [military] | Organizations | %<>*<> [organization/name] AND NOT %
 [military] | Money | %<>*<> [money]
 [organization/name] | People | %<>*<> [person/name]
 [organization/name] | Places | %<>*<> [location]
 [organization/name] | Organizations | %<>*<> [organization/name] AND NOT %
 [organization/name] | Money | %<>*<> [money]



Example Relationship Structure Guided Navigation Tips

Appendix B

```

<?xml version="1.0" encoding="utf-8" ?>
  <!DOCTYPE tiprules (View Source for full doctype...)>
= <tiprules>
- <!--
+++ Tips for Character +++
-->
  <rule1 desc="Character" source="[Person/Name]"
    target="[Person/Name]" score="5" />
  <rule1 desc="Place" source="[Person/Name]"
    target="[Location/Name]" score="5" />
  <rule1 desc="Organization" source="[Person/Name]"
    target="[Organization/Name]" score="5" />
= <rule desc="Action" includesDeepTips="true" score="5">
  <template source="[Person/Name]" action="$NONE"
    target="$NONE" />
  <query source="$IQLSource AND [Person/Name]" biDir="true"
    />
  </rule>
- <!--
+++ Tips for Place +++
-->
  <rule1 desc="Place" source="[Location/Name]"
    target="[Location/Name]" score="5" />
  <rule1 desc="Character" source="[Location/Name]"
    target="[Person/Name]" score="5" />
  <rule1 desc="Organization" source="[Location/Name]"
    target="[Organization/Name]" score="5" />
= <rule desc="Action" includesDeepTips="true" score="5">
  <template source="[Location/Name]" action="$NONE"
    target="$NONE" />
  <query source="$IQLSource AND [Location/Name]"
    biDir="true" />
  </rule>
- <!--
+++ Tips for Organization +++
-->
  <rule1 desc="Organization" source="[Organization/Name]"
    target="[Organization/Name]" score="5" />
  <rule1 desc="Character" source="[Organization/Name]"
    target="[Person/Name]" score="5" />
  <rule1 desc="Place" source="[Organization/Name]"
    target="[Location/Name]" score="5" />
= <rule desc="Action" includesDeepTips="true" score="5">
  <template source="[Organization/Name]" action="$NONE"
    target="$NONE" />
  <query source="$IQLSource AND [Organization/Name]"
    biDir="true" />
  </rule>
- <!--
+++ Tips for LocationRel +++
-->

```

```

<rule1 desc="Place" source="[LocationRel]"
  target="[Location/Name] AND NOT $IQLSource" score="5" />
<rule1 desc="Character" source="[LocationRel]"
  target="[Person/Name]" score="5" />
<rule1 desc="Organization" source="[LocationRel]"
  target="[Organization/Name]" score="5" />
= <rule desc="Action" includesDeepTips="true" score="5">
  <template source="[LocationRel]" action="$NONE"
    target="$NONE" />
  <query source="$IQLSource AND [LocationRel]" biDir="true" />
</rule>
- <!--
+++ Country: $COUNTRY>*>$EXISTS +++
-->
= <rule desc="Country" includesDeepTips="true" score="5">
  <template source="$COUNTRY" target="$EXISTS" />
  <query source="[Country]" action="$IQLAction"
    target="$IQLTarget" biDir="false" />
</rule>
- <!--
+++ People: $PERSON>*>$EXISTS +++
-->
= <rule desc="People" includesDeepTips="true" score="5">
  <template source="$PERSON" target="$EXISTS" />
  <query source="[Person/Name]" action="$IQLAction"
    target="$IQLTarget" biDir="false" />
</rule>
- <!--
+++ Money: $MONEY>*>$EXISTS +++
-->
= <rule desc="Money" includesDeepTips="true" score="5">
  <template source="$MONEY" target="$EXISTS" />
  <query source="$IQLTarget" action="$IQLAction"
    target="[Money]" biDir="false" />
</rule>
- <!--
+++ Places: $LOCATION>*>$EXISTS +++
-->
= <rule desc="Places" includesDeepTips="true" score="5">
  <template source="$LOCATION" target="$EXISTS" />
  <query source="$IQLTarget" action="$IQLAction"
    target="[location/name]" biDir="false" />
</rule>
- <!--
+++ Time: $TIME>*>$EXISTS +++
-->
= <rule desc="Time" includesDeepTips="true" score="5">
  <template source="$TIME" target="$EXISTS" />
  <query source="$IQLTarget" action="$IQLAction"
    target="[Temporal]" biDir="false" />
</rule>
- <!--

```

```

+++ Terrorist activities associated with $IQLSource: $TERROR_STATES>*>* +++
-->
= <rule desc="Terrorist activities associated with $IQLSource"
  includesDeepTips="true" score="5">
  <template source="$TERROR_STATES" />
  <query source="$IQLSource" target="[location] and not
    $IQLSource" context="terrorist or bomb or suicide"
    biDir="true" />
  </rule>
- <!--
+++ The $IQLSource War: $WAR_STATES>*>* +++
-->
= <rule desc="The $IQLSource War" includesDeepTips="true"
  score="5">
  <template source="$WAR_STATES" />
  <query source="$IQLSource war" target="[name] and not
    $IQLSource" biDir="true" />
  </rule>
- <!--
+++ $IQLSource and Cold War: $COMMIES_STATES>*>* +++
-->
= <rule desc="$IQLSource and Cold War" includesDeepTips="true"
  score="5">
  <template source="$COMMIES_STATES" />
  <query source="$IQLSource" target="[entity] and not
    $IQLSource" context="cold war" biDir="true" />
  </rule>
- <!--
+++ Hostilities between $IQLSource and $IQLTarget: [Country]>*>[Country] +++
-->
= <rule desc="Hostilities between $IQLSource and $IQLTarget"
  includesDeepTips="true" score="5">
  <template source="[Country]" target="[Country]" />
  <query source="$IQLSource" action="(invade or occupy or
    attack or raid or bomb or destroy or fight or assail or
    assault or combat or shell or beat or crush or blast or fire
    or enter or shoot or win or lose or take over)"
    target="$IQLTarget;war" biDir="true" />
  </rule>
- <!--
+++ $IQLSource in $IQLTarget: [Person]>*>[Location] +++
-->
= <rule desc="$IQLSource in $IQLTarget" includesDeepTips="true"
  score="5">
  <template source="[Person]" target="[Location]" />
  <query source="$IQLSource" action="visit or live"
    target="$IQLTarget" biDir="false" />
  </rule>
- <!--
+++ $IQLSource relates to $IQLTarget: $EXISTS>*>$EXISTS +++
-->

```

```

= <rule desc="$IQLSource relates to $IQLTarget"
  includesDeepTips="true" score="5">
  <template source="$EXISTS" target="$EXISTS" />
  <query source="$IQLSource" target="$IQLTarget" biDir="true"
    />
  </rule>
- <!--
+++ $IQLSource <> $IQLAction : $EXISTS> $EXISTS > $NONE +++
-->
= <rule desc="$IQLSource <> $IQLAction" includesDeepTips="true"
  score="5">
  <template source="$EXISTS" action="$EXISTS" target="$NONE"
    />
  <query source="$IQLSource" action="$IQLAction"
    target="$IQLTarget" biDir="true" />
  </rule>
- <!--
+++ Place: $EXISTS>*>$EXISTS +++
-->
= <rule desc="Place" includesDeepTips="true" score="5">
  <template source="$EXISTS" target="$EXISTS" />
  <query source="[Location/Name] and not $IQLSource and not
    $IQLTarget" context="$IQLSource $IQLTarget" biDir="true"
    />
  </rule>
- <!--
+++ Character: $EXISTS>*>$EXISTS +++
-->
= <rule desc="Character" includesDeepTips="true" score="5">
  <template source="$EXISTS" target="$EXISTS" />
  <query source="[Person/Name] and not $IQLSource and not
    $IQLTarget" context="$IQLSource $IQLTarget" biDir="true"
    />
  </rule>
- <!--
+++ Organization: $EXISTS>*>$EXISTS +++
-->
= <rule desc="Organization" includesDeepTips="true" score="5">
  <template source="$EXISTS" target="$EXISTS" />
  <query source="[Organization/Name] and not $IQLSource and
    not $IQLTarget" context="$IQLSource $IQLTarget"
    biDir="true" />
  </rule>
- <!--
+++ Action: $EXISTS>*>$EXISTS +++
-->
= <rule desc="Action" includesDeepTips="true" score="5">
  <template source="$EXISTS" target="$EXISTS" />
  <query source="$IQLSource" target="$IQLTarget" biDir="true"
    />
  </rule>
- <!--

```

```

+++ Place: $EXISTS>$EXISTS>$NONE +++
-->
= <rule desc="Place" includesDeepTips="true" score="5">
  <template source="$EXISTS" action="$EXISTS" target="$NONE"
    />
  <query source="[Location/Name] and not $IQLSource"
    action="$IQLAction" context="$IQLSource" biDir="true" />
</rule>
- <!--
+++ Character: $EXISTS>$EXISTS>$NONE +++
-->
= <rule desc="Character" includesDeepTips="true" score="5">
  <template source="$EXISTS" action="$EXISTS" target="$NONE"
    />
  <query source="[Person/Name] and not $IQLSource"
    action="$IQLAction" context="$IQLSource" biDir="true" />
</rule>
- <!--
+++ Organization: $EXISTS>$EXISTS>$NONE +++
-->
= <rule desc="Organization" includesDeepTips="true" score="5">
  <template source="$EXISTS" action="$EXISTS" target="$NONE"
    />
  <query source="[Organization/Name] and not $IQLSource"
    action="$IQLAction" context="$IQLSource" biDir="true" />
</rule>
- <!--
+++ Substitution variables +++
-->
= <var name="CITY">
  <term>city</term>
  <term>town</term>
  <term>village</term>
  <term>borough</term>
</var>
= <var name="COMMIES_STATES">
  <term>Russia</term>
  <term>Soviet</term>
  <term>China</term>
  <term>Vietnam</term>
  <term>North Korea</term>
  <term>Cuba</term>
</var>
= <var name="COUNTRY">
  <term>country</term>
  <term>state</term>
  <term>nation</term>
</var>
= <var name="LOCATION">
  <term>location</term>
  <term>place</term>
</var>

```

```

= <var name="PERSON">
  <term>person</term>
  <term>people</term>
  <term>someone</term>
  <term>anyone</term>
</var>
= <var name="TERROR_STATES">
  <term>iraq</term>
  <term>iran</term>
  <term>syria</term>
  <term>north korea</term>
  <term>libya</term>
</var>
= <var name="TIME">
  <term>time</term>
  <term>date</term>
</var>
= <var name="WARRING">
  <term>invade</term>
  <term>occupy</term>
  <term>attack</term>
  <term>raid</term>
  <term>bomb</term>
  <term>destroy</term>
  <term>fight</term>
  <term>assail</term>
  <term>assault</term>
  <term>combat</term>
  <term>shell</term>
  <term>beat</term>
  <term>crush</term>
  <term>blast</term>
  <term>fire</term>
  <term>enter</term>
  <term>shoot</term>
  <term>win</term>
  <term>lose</term>
  <term>take over</term>
</var>
= <var name="WAR_STATES">
  <term>Iraq</term>
  <term>Iran</term>
  <term>vietnam</term>
  <term>Korean</term>
</var>
</tiprules>

```



InFact® Ontology Design

Appendix C

1.0 Introduction

This document addresses incorporating domain-specific, customer-specified ontologies.

The Ontology is a component of the InFact system that allows an administrator to alter the way sentences are parsed. The Ontology also includes Thesaurus abilities including the ability to add synonyms, and acronyms. In addition, the ontology allows tagging words with different categorical paths that may be used to filter query results, and thus provide more accurate information retrieval.

The Ontology is a component of the InFact system, which consists of the following primary functionality:

- Ability to add noun phrases to InFact
- Ability to add acronyms, synonyms to InFact
- Ability to add 1 or more categorical paths to an Ontology term
- Ability to support very large (> 1 million terms) Ontologies.

An ontology main goal is to improve efficacy during search. The ontology addresses this goal by helping the parser recognize the correct part of speech of specialized (corpus dependent) vocabulary terms such as complex noun phrases, as well as by tagging the sentence terms with special corpus dependent semantic attributes, such as Ontology Paths. An ontology also allows user to specify synonyms. These are to provide extremely useful filters during search.

Before parsing a sentence in NLP, ontology terms (noun phrases) are replaced by internal named entities. After the NLP parsing the same named entities are replaced with the original noun phrases. This is to avoid ambiguities with the NLP parser ontologies, lexical, syntactic, and semantic rules. This substitution mechanism is employed during both indexing and searching.

Ontology items are preferably defined in human readable and editable XML files.

Ontology items can optionally add (in addition to noun phrase/semantic attribute pair) a synonym for the ontology term. The synonym is then to normalize the ontology term during internal named entity to ontology term re-substitution (after NLP parsing).

2.0 Ontology philosophy and Data Layout

The ontology consists of Entities, which contain Terms. Each Term within an entity is considered to be a synonym for all terms within this Term. The Entities exist independently and they are not merged. In one embodiment the synonym relationship is not transitive. The definition of synonym according to Webster <http://www.m-w.com> is: one of two or more words or expressions of the same language that have the same or nearly the same meaning in some or all senses. This definition implies grouping, not transitivity.

Example:

WHO is a synonym for World Health Organization. WHO is also a synonym for White House Office, but White House Office is not a synonym for World Health Organization (unless it has been specified as such). Search for WHO should also search for White House and World Health Organization, but search for World Health Organization should search only for WHO and not White House Office.

The ignoreCase and doMorph flags are typically set independently on ontology, Entity and term levels.

Example:

The Entity that contains: US, USA, United States, and United States of America should have ignoreCase set to yes, with the exception of term US, which should have ignoreCase set to no.

The Ontology Paths are typically set on Entity level. If some synonym(s) should have separate OPS, a new Entity should be created.

The ontology may specify optional normalization for Entities. Such normalization may be used for visualization and dataframe export purposes.

Example:

The entity that contains: US, USA, United States, and United States of America may be normalized to USA.

The DTD for the ontology can be defined as:

```
<!ELEMENT ontology ( entity+ ) >
<!ATTLIST ontology ignorecase NMTOKEN #REQUIRED >
<!ATTLIST ontology morph NMTOKEN #REQUIRED >

<!ELEMENT entity ( term+, path* ) >
<!ATTLIST entity ignorecase NMTOKEN #IMPLIED >
<!ATTLIST entity morph NMTOKEN #IMPLIED >

<!ELEMENT path EMPTY >
<!ATTLIST path value CDATA #REQUIRED >

<!ELEMENT term EMPTY >
<!ATTLIST term ignorecase NMTOKEN #IMPLIED >
<!ATTLIST term morph NMTOKEN #IMPLIED >
<!ATTLIST term normal NMTOKEN #IMPLIED >
<!ATTLIST term phrase CDATA #REQUIRED >
```

3.0 KWS Schema

Ontology Entity data is stored in a KWS appropriate manner in an ontology noun phrase index. A KWS "Document" exists for every Ontology Entity in the Ontology. A given document has the following KWS fields:

Field Name	Type	Description
entityId	Keyword - indexed stored	EntityID
flags	UnIndexed - stored	Flags for the entity (dimorph, ignoreCase, isNormal, useSubstitution, isStandard, number of words)
pathName	Keyword - indexed stored	Ontology Path with spaces replaced by '_'

Ontology Term data is stored in a KWS appropriate manner in an ontology noun phrase index. A KWS "Document" exists for every Ontology Term in the Ontology. A given document has the following KWS fields:

Field Name	Type	Description
------------	------	-------------

entityRefId	Keyword - indexed stored	EntityID to which term belongs to
text	UnIndexed - stored	Ontology Term
contents	Text – tokenized, Indexed, stored	Ontology Term or lowercase Ontology Term if ignoreCase is set to true
lccontents	UnStored - tokenized, Indexed	lowercase contents
concat	UnStored - tokenized, Indexed	Ontology Term or lowercase Ontology Term if ignoreCase is set to true with all spaces replaced by '_'
lcconcat	UnStored - tokenized, Indexed	Lowercase concat
phrase	Keyword - indexed stored	Ontology Term or lowercase Ontology Term if ignoreCase is set to true
lcphrase	Keyword - indexed stored	Lowercase phrase
flags	UnIndexed - stored	Flags for the term (dimorph, ignoreCase, isNormal, useSubstitution, isStandard, number of words)
pathName	Keyword - indexed stored	Ontology Path with spaces replaced by '_'

Each Ontology Path has spaces replaced by '_'. Ontology Path data is stored in a KWS appropriate manner in an ontology path index. A KWS "Document" exists for every Ontology Path in the Ontology. A given document has the following KWS fields:

Field Name	Type	Description
path	UnStored - tokenized indexed	Tokenized lowercase path
name	Keyword - indexed stored	Untokenized path
subpath_0 ... subpath_n	Text – tokenized, indexed, stored	Lowercase path substrings with _____ attached subpath_0 = if _____ subpath_1 = if/entity_____ subpath_2 = if/entity/person_____
osubpath_0 ... osubpath_n	UnIndexed - stored	path substrings subpath_0 = IF subpath_1 = IF/Entity subpath_2 = IF/Entity/Person
code	Keyword - Indexed, stored	Integer code

In addition to the paths, a Maximum Path Length is also stored in a KWS index as a separate KWS document with a "system" Keyword field set to "info", and UnIndexed "maxPathDepth" field containing a maximum path depth.

4.0 Example of an Ontology File

```
<?xml version="1.0" encoding="utf-8"?>
```

```
<!ELEMENT ontology ( entity+ ) >
<!ATTLIST ontology ignorecase NMTOKEN #REQUIRED >
<!ATTLIST ontology morph NMTOKEN #REQUIRED >
```

```
<!ELEMENT entity ( term+, path* ) >
<!ATTLIST entity ignorecase NMTOKEN #IMPLIED >
```

<!ATTLIST entity morph NMTOKEN #IMPLIED >

<!ELEMENT path EMPTY >

<!ATTLIST path value CDATA #REQUIRED >

<!ELEMENT term EMPTY >

<!ATTLIST term ignorecase NMTOKEN #IMPLIED >

<!ATTLIST term morph NMTOKEN #IMPLIED >

<!ATTLIST term normal NMTOKEN #IMPLIED >

<!ATTLIST term phrase CDATA #REQUIRED >

<ontology morph="yes" ignorecase="yes">

<entity morph="yes" ignorecase="yes">

<term phrase="USA" normal="yes" />

<term phrase="American" />

<term phrase="US" ignorecase="no" />

<term phrase="USA" />

<term phrase="U.S." />

<term phrase="U.S.A." />

<term phrase="United States" />

<term phrase="United States of America" />

<path value="IF/Entity/Location/Country" />

</entity>

<entity>

<term phrase="World Health Organization" morph="yes" ignorecase="yes" />

<term phrase="WHO" morph="no" ignorecase="no" />

<path value="IF/Entity/Organization/Name" />

</ entity >

<entity morph="no">

<term phrase="White House Office" ignorecase="yes" />

<term phrase="WHO" ignorecase="no" />

<path value="IF/Entity/Organization/Government" />

<path value="IF/Entity/Organization/Name" />

</entity>

<entity >

<term phrase="United Airlines" normal="yes" />

<term phrase="UAL" />

<path value="IF/Entity/Organization/Name" />

</entity>

<entity morph="yes">

<term phrase="Second World War" normal="yes" />

<term phrase="WW2" />

<term phrase="WWII" />

<term phrase="2nd World War" />

<term phrase="World War II" />

<path value="IF/Entity/Temporal/Time_Period" />

<path value="IF/Entity/Temporal/Event" />

</entity>

<entity ignorecase="yes">

<term phrase="Franklin D. Roosevelt" normal="yes" />

<term phrase="FDR" />

<term phrase="F.D.R." />

<term phrase="F.D.R " />

```
<term phrase="Franklin D Roosevelt " />
<path value="IF/Entity/Person/Male" />
<path value="IF/Entity/Person/Name" />
</entity>
```

```
<entity morph="yes" ignorecase="yes">
  <term phrase="Al Queda " normal="yes" />
  <term phrase="Al Qaeda" />
  <term phrase="Al Qaida" />
  <term phrase="Al Qa'eda" />
  <term phrase="Al-Qa'eda" />
  <term phrase="Al-Queda" />
  <term phrase="Al Qu'eda" />
  <term phrase="Al-Qu'eda" />
  <term phrase="Al Qa'ida" />
  <term phrase="Al Qa'idah" />
  <term phrase="Al-Qa'ida" />
  <term phrase="Al-Qa'idah" />
  <term phrase="Al-Qida" />
  <term phrase="Al-Quaeda" />
  <path value="IF/Entity/Organization/Military" />
  <path value="IF/Entity/Organization/Name" />
</entity>
```

```
</ontology>
```

CLAIMS

1. A method in a computing system for performing enhanced electronic relationship searches on a corpus of documents, comprising:

receiving a designated query indicating one or more terms;

in response to the received designated query, determining a plurality of matching rules that correspond to at least one indicated term of the designated query, each rule including a term used to decide whether the designated query matches the rule and including a relationship search specification string that specifies a corresponding relationship search to be executed as a search against the corpus when the rule is matched, the relationship search indicating one or more terms and associated syntactic and/or semantic information used to convey how the one or more terms are used to be understood in relation to the corpus;

for each of the plurality of determined matching rules, determining and executing the corresponding relationship search specified by the relationship search specification string in parallel against the corpus to validate whether the corresponding relationship search produces a successful result, thereby causing multiple different relationship searches to be executed against the same corpus in response to a same designated query at the same time; and

for each of a plurality of the executed corresponding relationship searches that have been validated as producing a successful result, forwarding an indication of the corresponding relationship search as a suggested search tip, each suggested search tip indicative of a different relationship search applicable to the same designated query.

2. The method of claim 1 wherein the receiving the designated query further comprises receiving an indication of one or more components of a relationship query.

3. The method of claim 1 or 2 wherein the receiving the designated query further comprises receiving an indication of one or more keywords.

4. The method of at least one of claims 1 to 3 wherein the receiving the designated query comprises receiving an indication of an entity or an entity type.

5. The method of claim 4 wherein the entity or the entity type is a source for an action or for an action type.

6. The method of at least one of claims 4 or 5 wherein the entity or the entity type is a target for an action or for an action type.

7. The method of at least one of claims 1 to 6 wherein the receiving the designated query comprises receiving an indication of an action or an action type.

8. The method of at least one of claims 1 to 7, further comprising:
for the one or more corresponding executed relationship searches that have been validated as producing a successful result,
evaluating results of the corresponding executed relationship search that produced the successful result to determine a plurality of additional specific relationship queries to serve as additional suggested searches; and

forwarding indications of the plurality of additional specific relationship queries as additional suggested search tips.

9. The method of at least one of claims 1 to 8 wherein, for the one or more corresponding executed relationship searches that have been validated as producing a successful result, the forwarding the indication of the corresponding relationship search as a suggested search tip further comprises:

evaluating results of the corresponding executed relationship search to determine whether the results meet a set of criteria for presenting as a suggested search tip; and

when it is determined that the results meet the criteria, forwarding an indication of the corresponding relationship search as a suggested search tip.

10. The method of at least one of claims 1 to 9 wherein a corresponding executed relationship search that produces a successful result is determined by whether search results are produced from executing the corresponding relationship search against the corpus.

11. The method of at least one of claims 1 to 10 wherein the suggested search tip is presented as a navigation tip in a visual user interface.

12. The method of at least one of claims 1 to 11, further comprising:
caching search results of one or more corresponding executed relationship searches that have been validated as producing a successful result; and
when a suggested searched tip is indicated for execution, returning the cached results of the relationship search that corresponds to the suggested search tip.

13. The method of at least one of claims 1 to 12 wherein the rules are specified by a configuration file, and further comprising:
parsing the configuration file to determine, for each rule, a relationship search specification string that specifies a corresponding relationship query; and
for each determined relationship search specification string, storing the corresponding relationship query along with information for matching the corresponding relationship query to a received designated query.

14. The method of at least one of claims 1 to 13 wherein the rules are specified using a tag-based scripting language.

15. The method of at least one of claims 1 to 14, further comprising presenting indications of the corresponding relationship searches in a ranked order.

16. The method of at least one of claims 1 to 15 wherein the executing the corresponding relationship searches is performed in parallel.

17. The method of at least one of claims 1 to 16 wherein the determining the plurality of matching rules that correspond to the designated query further comprises:
determining a set of matching ontology paths that correspond to an indicated entity or entity type of the designated query.

18. The method of at least one of claims 1 to 17, further comprising:
when a suggested search tip is indicated for execution, executing the corresponding relationship search; and
presenting search results of the executed relationship search with indications of the components of the relationship query performed by the executed relationship search.
19. The method of claim 18 wherein the indications of the components in the presented results are used to facilitate teaching syntax of a specific query language.
20. A computer-readable medium containing instructions that, when executed, controls a computer processor to carry out the method of at least one of claims 1 to 19.
21. The computer readable medium of claim 20 comprising a computer memory.
22. The computer-readable medium of at least one of claims 20 or 21 wherein the receiving the designated query further comprises receiving an indication of one or more components of a relationship query.
23. The computer-readable medium of claim 22 wherein the one or more components of the relationship query are indicated using IQL.
24. The computer-readable medium of at least one of claims 20 to 23 wherein the receiving the designated query further comprises receiving an indication of more or more keywords.
25. The computer-readable medium of at least one of claims 20 to 24 wherein the receiving the designated query comprises receiving an indication of one or more of an entity, an entity type, an action, or an action type.

26. The computer-readable medium of at least one of claims 20 to 25, the method further comprising:

for the one or more executed corresponding relationship searches that have been validated as producing a successful result,

evaluating results of the executed corresponding relationship search that produced the successful result to determine a plurality of more specific relationship queries to serve as additional suggested searches; and

forwarding indications of the plurality of more specific relationship queries as additional suggested search tips.

27. The computer-readable medium of at least one of claims 20 to 26, the method further comprising:

caching search results of the one or more executed corresponding relationship searches that have been validated as producing a successful result; and

when a suggested searched tip is indicated for execution, returning the cached results of the executed relationship search that corresponds to the suggested search.

28. The computer-readable medium of at least one of claims 20 to 27 wherein the rules are specified in a file using a tag-based scripting language.

29. The computer-readable medium of at least one of claims 20 to 28, the method further comprising presenting the one or more suggested search tips in a ranked order.

30. The computer-readable medium of at least one of claims 20 to 29 wherein the executing of the corresponding relationship searches is performed in parallel.

31. A relationship search navigation tip system comprising:

a computer processor;

a data repository configured to maintain rules that map between one or more relationship query components and corresponding relationship query specification strings that

specify corresponding relationship queries to be executed as searches against a corpus when the corresponding rule is matched;

a storage medium comprising a tip search manager that, when executed on the computer processor, is configured to parse the rules stored in the data repository to generate rule mappings between potential input queries and corresponding relationship searches to be executed, each relationship search indicating one or more terms and associated syntactic and/or semantic information used to convey how the one or more terms are used to be understood in relation to the corpus; and

a storage medium comprising a tip searcher that, when executed on the computer processor, is configured to receive an input query, determine, using the tip search manager, a plurality of rule mappings that correspond to the input query by matching a term of the input query to a plurality of the stored rules, execute the plurality of different relationship search queries that are specified by the relationship search specification strings of the determined plurality of rule mappings to generate corresponding search results, and return indications to the executed relationship searches that generated successful corresponding search results.

32. The navigation tip system of claim 31 wherein the indications to the executed relationship searches that generated successful corresponding search results are visual navigation tips.

33. The navigation tip system of at least one of claims 31 or 32 wherein the tip searcher is further configured to provide additional indications to relationship searches that correspond to specific entities or specific actions discovered during the executed plurality of relationship searches.

34. The navigation tip system of claim 33 wherein the additional indications comprise deep tips.

35. The navigation tip system of at least one of claims 31 to 34 wherein the tip searcher is configured to execute the relationship searches that correspond to the determined plurality of mappings in parallel with each other.

36. The navigation tip system of at least one of claims 31 to 35 wherein the tip searcher is configured to execute the relationship searches that correspond to the determined plurality of mappings asynchronously.

37. A computer-implemented method for performing a relationship search against a corpus of documents comprising:

receiving from a user a relationship query indicating one or more terms and semantic and/or syntactic information associated with the one or more terms, the received relationship query including a separate document constraint portion, the terms and associated semantic and/or syntactic information specifying what relationships in one or more documents of the corpus will match the query using, at least in part, the indicated semantic and/or syntactic information, the document constraint portion specifying a document constraint operator symbol or phrase and a designated ontology path specification;

executing a relationship search against the corpus of documents for a designated entity and/or designated action of the received relationship query to yield a plurality of indications of documents in the corpus, where each indicated document contains at least one relationship involving the designated entity and/or designated action as specified by the indicated terms and associated semantic and/or syntactic information indicated by the query;

as a filter of results obtained by executing the relationship search, determining whether each of the indicated documents also contains a term that matches the designated ontology path specification of the document constraint portion of the received indicated relationship query; and

for each determined document that satisfies the document constraint portion, returning an indication of the document as a successful result of the relationship search.

38. A computer-implemented method for performing a relationship search against a corpus of documents comprising:

receiving from a user a designated relationship query indicating one or more terms and semantic and/or syntactic information associated with the one or more terms, the received designated relationship query including a separate context operator portion, the terms and associated semantic and/or syntactic information specifying what relationships in one or more documents of the corpus will match the query using, at least in part, the indicated semantic and/or syntactic information, the context portion specifying a context operator symbol or phrase and an entity and/or an action;

executing a relationship search that corresponds to the designated relationship search query to yield one or more indications of documents in the corpus, where each document contains at least one sentence or clause that matches the relationship specified by the indicated terms and associated semantic and/or syntactic information of the designated relationship query; and

for each indicated document, determining whether a context surrounding the at least one matching sentence or clause, within a determined number of sentences or clauses before or after the at least one matching sentence or clause, further matches the entity or the action specified by the context operator portion of the designated relationship query, wherein the context is a sentence or a clause that is distinct from the at least one matching sentence or clause.

39. The method of claim 38 wherein the context includes a sentence before the at least one matching sentence.

40. The method of at least one of claims 38 or 39 wherein the context includes a sentence after the at least one matching sentence.

41. The method of at least one of claims 38 to 40 wherein the designated query specifies at least one of a source entity, a target entity, or an action and specifies a context operator along with a corresponding entity or action.






Fig. 2





Fig. 3

4/31

SQE Processing
Overview**Fig. 4**

5/31

globalSecurity.org - Search - Reliable Security Information - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://globalSecurity.org/search/jspr/s.jsp?Search=exec& Google

Best of the Web Channel Guide Customize Links Free Hotmail Internet Start Liquid Music Network Microsoft RealPlayer

SEARCH GLOBAL SECURITY.ORG

Powered by **Infact**

Search | History | Help | Contact

Try your own Fact Search

Search

Fact Search results 1 - 100 of about 313; 501 Page 1 of 4 Next

View Report Retrieve by date Sort page by: Action Similarity Sort

Source	Action	Target
Bush Appoints Bolton : to UN Post Without Senate	01 Aug 2005 : vote	VOA 502
Mrs. A Bush	vote	-
George W Bush	vote	-
Vice : President Bush	vote : in November	-
George Bush	vote : to break	-
	select[2]	Anthony A. Principi : of California
President Bush	select[2]	Tom Ridge
President Bush	select[2]	-
President : George H W Bush	select : as site for special Thanksgiving Day meal with Sailor Marine thousand of mile from home	Nassau

Done

503

Fig. 5

6/31

GlobalSecurity.org - Search - Reliable Security Information - Media Alerts

File Edit View History Bookmarks Tools Help

Best of the Web Channel Guide Customize Links Free Hotmail Internet Start Liquid Music Network Microsoft RealPlayer

SEARCH GLOBAL SECURITY.ORG

Powered by **Infact**

Search | History | Help | Contact

Try your own Fact Search

bush > vote > * Search

Fact Search results 1 - 100 of about 297: 604

View Report Retrieve by relevance

Date	Source	Action
11/09/2006	Bush	select
11/09/2006	President Bush	choose : as Donald Rumsfeld successor
11/09/2006	09 November 2006 in-depth Coverage : President Bush	choose : By Greg Flakus Houston
09/27/2006	Bush administration	choose/31
08/20/2006	Bush : administration	choose
06/29/2006	President Bush	choose
06/21/2006	George H W Bush : decision	select : as secretary of defense
06/13/2006	Bush administration	Thursday : vote : on

603

601

602

Sort page by: Date Source Action Target Action Frequency Action Similarity

page 1 of 3 Next

600 existing : venue military courts-martial civilian : court Donald Rumsfeld

Fig. 6

7/31

terrorist > attack > london

Search

Fact Search results 1 - 4:

Click this button to export or view an HTML report of results, or customize a tab-delimited file for export.

View Report Go

703 704

Sort page by: Action Similarity Sort

Source	Action	Target
terrorist	first : attack : on 7 July then again on 21 July	London transport : system
21/22/23/24/25/26 : UTC Dates : 01 possible terrorist	attack	US Team River : London
terrorist : bombing bus : on Thursday	hit	subway train
terrorist	terrorist	three : London
		London

Fig. 7

8/31




Fig. 8

9/31

Infact

search | Corpus | Preferences | History | Help | Contact

Try your own Fact Search

[organization/name]>buy>[organization/name]^[money] ~computer

Search

Fact Search results 1 - 55:

901 902 903

Export Data Frame Go

Nested search: Set

Retrieves by date

Sort page by: Date Publish Sort

Date Published	Source	Action	Target
08/08/1997	Hikari Tsushin Inc	buy : for three billion yen	Sofmap Co Ltd : stock
07/28/1997	Intel	acquire	Acron Technology : for \$17.50-\$3.50 above Friday Nasdaq closing price
07/25/1997	Chartwell Group Inc	buy	Chip
07/21/1997	ThermoTrex Corp	acquire	200,000 additional : share : of Computer Integration Corp stock for \$1.13
07/18/1997	Sun	buy : for \$185 million	Computer Communications Specialists Inc : for \$10 million to post-voting acquisition
			Encore Computer Corp. storage business : asset

Fig. 9

10/31

1001

Step 1: Specify columns

Column	Column data
1	Source1002 Remove
2	Action1003 Remove
3	Target1004 Remove
4	Prepositional Phrase1005 Remove
5	Date Published1006 Remove
6	Creator Country1007 Remove

Preview (tab delimited):1009

Mikari Teushin Inc	buy	Softap Co Ltd	3.05E9 yen	Fri Aug 08 00:00:00 PDT 1997	JAPAN
INTEL CORP	acquire	Nasdaq	\$3.50	Mon Jul 28 00:00:00 PDT 1997	USA
INTEL CORP	acquire	Nasdaq	\$17.50	Mon Jul 28 00:00:00 PDT 1997	USA
Chartwell Group Inc	buy	Computer Integration Corp	\$1.13	Fri Jul 25 00:00:00 PDT 1997	USA
ThermoFlex Corp	acquire	Computer Communications Specialists Inc	\$1.85E8	Fri Jul 18 00:00:00 PDT 1997	USA
SUNOCO INC	buy	Encore Computer Corp.	\$1.85E8	Fri Jul 18 00:00:00 PDT 1997	USA
SUNOCO INC	acquire	Encore Computer Corp.	\$1.85E8	Fri Jul 18 00:00:00 PDT 1997	USA
SUNOCO INC	acquire	Encore Computer Corp.	\$1.85E8	Fri Jul 18 00:00:00 PDT 1997	USA
SUNOCO INC	acquire	Encore Computer Corp	\$1.85E8	Fri Jul 18 00:00:00 PDT 1997	USA
SUN MICROSYSTEMS INC	buy	Encore Computer Corp	\$1.85E8	Thu Jul 17 00:00:00 PDT 1997	USA

Add Column

Export Data

1008

1010

Fig. 10

1101

Fig. 11

Fig. 11

12/31




Fig. 12

13/31




Fig. 13

14/31




Fig. 14



16/31



1604

17/31




Fig. 17A

18/31



Fig. 17B

19/31

Infact Book Search - Search, explore and download free ebooks online - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

Customize Links Free HTML ReaPlayer Windows Marketplace Windows Media Windows

Google

Top 10 FAQ About

United Kingdom

Powered by **Infact** Book Search Engine

From All Books

Results Page 1 for United Kingdom related to Any Religion ~ 1810

1801 ~ 1802

Confessions of a Young Man - Page 23 ~ 1820a ~ England ~ 1820b ~ Protestantism ~ 1820c ~ 1820

1820a ~ England ~ 1820b ~ Protestantism ~ 1820c ~ 1820

...of this level will be found on the sinister side of the title-page. THE KNIGHTS OF THE THISTLE. The most ancient order of the Thistle was founded by James V. of Scotland, 1540, and revived by James II., king of Great Britain, 1687, incorporated by Queen Anne, whose statutes were confirmed by George I. The order consists of the sovereign and twelve brethren or knights. Their motto is the national motto...

The Manual of Heraldry; Fifth Edition - Page 31

king / Great Britain / James II ~ revive ~ order of the Thistle ~ 1821

The Conquest of New France - Page 8

struggle / Roman Catholic Protestant ~ bring ~ revolution / England / Europe / America ~ 1822

...of this level will be found on the sinister side of the title-page. THE KNIGHTS OF THE THISTLE. The most ancient order of the Thistle was founded by James V. of Scotland, 1540, and revived by James II., king of Great Britain, 1687, incorporated by Queen Anne, whose statutes were confirmed by George I. The order consists of the sovereign and twelve brethren or knights. Their motto is the national motto...

United Kingdom related to

places

Any Place

France

United States

Spain

Germany

Ireland

organizations

Any Organization

navy


parliament

House of Commons

House of Lords

Done

Fig. 18A



21/31

Infact BookSearch - Search, explore and download free books online - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

Customize Links [L] Free Hotmail [H] RealPlayer [R] Windows Marketplace [W] Windows Media [M] Windows [W]

Google [G] Search [S] PageRank [P] Check [C] AutoLink [A] Options [O]

Top 10 FAO About

United Kingdom [U] Search Books [S]

Powered by **Infact**
Book Search Engine

From All Books

1902 1903

Results Page 1 for United Kingdom related to France

A School History of the United States - Page 29

1904

Great Britain → declare → war / France

...of the British vessel Trent, stopped her on the high seas, and took four passengers off her, and brought them prisoners to Fort Warren, in Boston harbor. Mason and Silldell are the two we remember, Confederate envoys to France and Great Britain. Over the whole North burst into glorious joy. Our Secretary of the Navy wrote to Wilkes his congratulations. Congress voted thanks to him, government...

A Straight Deal - Page 26

1905

Silldell / Mason → be → envoy / France / two / Great Britain

...of the British vessel Trent, stopped her on the high seas, and took four passengers off her, and brought them prisoners to Fort Warren, in Boston harbor. Mason and Silldell are the two we remember, Confederate envoys to France and Great Britain. Over the whole North burst into glorious joy. Our Secretary of the Navy wrote to Wilkes his congratulations. Congress voted thanks to him, government...

A History of Modern Europe, 1792-1878 - Page 201

1906

Great Britain → demand → assistance / Great Britain

...of the British vessel Trent, stopped her on the high seas, and took four passengers off her, and brought them prisoners to Fort Warren, in Boston harbor. Mason and Silldell are the two we remember, Confederate envoys to France and Great Britain. Over the whole North burst into glorious joy. Our Secretary of the Navy wrote to Wilkes his congratulations. Congress voted thanks to him, government...

United Kingdom related to characters

Any Character

Victoria

Baron

Earl

Louis

Mason

places

Any Place

France

United States

Spain

Germany

Done

Fig. 19




Fig. 20




Fig. 21

24/31

The screenshot displays the InFact Book Search website. The browser window shows the URL <http://books.infact.com/booksearch/jsp/bookoverview.jsp>. The page features a search bar with the text 'Steele' and a 'Search in this book' button. Below the search bar, the title 'The Lives of the Poets of Great Britain and Ireland : Cibber, Theophilus, 1703-1758' is displayed. The page is divided into two main sections: 'Results Page 1 for Steele as a character' and 'Page 20'. The 'Results Page 1' section includes a list of related characters (Any Character, Congreve, Doris, Christian Hero, John Dennis) and actions (Write, Recommend, Celebrate, Contract, Appoint). The 'Page 20' section includes a list of related characters (Any Character, Congreve, Doris, Christian Hero, John Dennis) and actions (Write, Recommend, Celebrate, Contract, Appoint). The page also features a 'download book' button and a 'Project Gutenberg' link.

Steele related to characters

- Any Character
- Congreve
- Doris
- Christian Hero
- John Dennis

actions

- Write
- Recommend
- Celebrate
- Contract
- Appoint

Results Page 1 for Steele as a character 2207

Page 20

...a song before she was seven years old. She is the author of fifteen plays; her talent is comedy, particularly the contrivance of the plots, and incidents. Sir Richard Steele, in one of his Tallers, speaking of the Busy Body, thus recommends it. The plot, and incidents of the play, are laid with that subtilty, and spirit, which is peculiar to females of wit, and is very seldom well performed by those...

Page 32

...of Homer, translated by our author, and the same passages by Pope and Tickell, in which comparison the palm is very deservedly yielded to Pope. Our author wrote a Satire called Doris, celebrated by Sir Richard Steele, who was a warm friend to Mr. Congreve. He also wrote the Judgment of Paris, a Masque; and the Opera of Semele; of these, the former was acted with great applause, and the latter is

download book

15 Project Gutenberg

Infact book index

- characters
- places
- organizations

United Kingdom

Top 10 **FAQ** **About**

Search Books

Steele **search in this book**

The Lives of the Poets of Great Britain and Ireland : Cibber, Theophilus, 1703-1758

Infact Book Search **Search, explore and download free books online - Mozilla Firefox**

File Edit View Go Bookmarks Tools Help

Customize Links **Free Hotmail** **RealPlayer** **Windows Marketplace** **Windows Media** **Windows**

Google **Search** **PageRank** **Check** **AdBlock** **AdBlock** **Options**

Powered by InFact **Book Search Demo**

Fig. 22

[illegible]

Fig. 23

26/31

Infact Book Search - Search, explore and download free ebooks online - Mozilla Firefox

File Edit View Go Bookmarks Tools Help

http://books.infact.com/booksearch/jsp/pagesearch.jsf

Customize Links: ☐ Free Hotmail ☐ RealPlayer ☐ Windows Marketplace ☐ Windows Media ☐ Windows

Google

Top 10 FAO About United Kingdom

Infact Book Search Demo

The Lives of the Poets of Great Britain and Ireland : Cibber, Theophilus, 1703-1758

Steele related to Any Character

Results Page 1 for Steele related to Any Character 2414

Page 32

Richard Steele → celebrate → Doris 2415

JOHN DENNIS → contract → B / Richard Steele 2416

infact book index

- characters
- places
- organizations

download book

Project Gutenberg

Steele related to Any Character

Steele related to

characters

Any Character 2413

Congreve

Doris

Christian Hero

John Dennis

actions

Wife

Recommend

Celebrate


Contract

Appoint


John and bad angels interlarded themselves in the action, and his hero, the Duke of Marlborough, enjoys a large share of the celestial protection. Mr. Dennis had once contracted a friendship[B] with Sir Richard Steele, whom he afterwards severely attacked. Sir Richard had promised that he would take some opportunity of mentioning the works in public with advantage, and endeavour to raise the


Fig. 24

27/31

**Fig. 25**

28/31

**Fig. 26**

29/31**Fig. 27**





Fig. 28

31/31

**Fig. 29**

From All Books

Results Page 1 for United Kingdom related to France



A School History of the United States - Page 29

Great Britain → declare → war / France

...of the British vessel Trent, stopped her on the high seas, and took four passengers off her, and brought them prisoners to Fort Warren, in Boston harbor. Mason and Slidell are the two we remember. Confederate envoys to France and Great Britain. Over this the whole North burst into glorious joy. Our Secretary of the Navy wrote to Wilkes his congratulations, Congress voted its thanks to him, governors...



A Straight Deal - Page 26

Slidell / Mason → be → envoy / France / two / Great Britain

...of the British vessel Trent, stopped her on the high seas, and took four passengers off her, and brought them prisoners to Fort Warren, in Boston harbor. Mason and Slidell are the two we remember. Confederate envoys to France and Great Britain. Over this the whole North burst into glorious joy. Our Secretary of the Navy wrote to Wilkes his congratulations, Congress voted its thanks to him, governors...



A History of Modern Europe, 1792-1873 - Page 201

Great Britain → demand → assistance / Great Britain

...month, a few more intensions and policy, and of importance as the power given to The Regency of Lisbon, justly treating these outrages as the act of the Spanish Government, and appealing to the treaties which bound Great Britain to defend Portugal against foreign attack,

1904

1905

1906

1901

United Kingdom related to

characters
Any Character
Victoria
Baron
Earl
Louis
Mason

places

Any Place
France
United States
Spain
Germany