




(86) Date de dépôt PCT/PCT Filing Date: 2011/09/21
(87) Date publication PCT/PCT Publication Date: 2012/03/29
(45) Date de délivrance/Issue Date: 2019/08/13
(85) Entrée phase nationale/National Entry: 2013/03/22
(86) N° demande PCT/PCT Application No.: US 2011/052586
(87) N° publication PCT/PCT Publication No.: 2012/040350
(30) Priorité/Priority: 2010/09/24 (US61/386,029)

(51) Cl.Int./Int.Cl. *G06F 17/27* (2006.01),
G06F 16/33 (2019.01)

(72) Inventeurs/Inventors:
FAN, JAMES J., US;
FERRUCCI, DAVID A., US;
GONDEK, DAVID C., US;
KALYANPUR, ADITYA A., US;
LALLY, ADAM P., US;
MURDOCK, JAMES W., IV, US;
ZADROZNY, WLODEK W., US;
LEVAS, ANTHONY, US

(73) Propriétaire/Owner:

(54) Titre : ESTIMATION ET APPLICATION DE CONFIANCE DE TYPE LEXICAL DE REPONSE
(54) Title: LEXICAL ANSWER TYPE CONFIDENCE ESTIMATION AND APPLICATION



(57) Abrégé/Abstract:

A system, method and computer program product for automatically estimating the confidence of a detected LAT to provide a more accurate overall score for an obtained candidate answer. A confidence "score" or value of each detected LAT is obtained, and the

(73) **Propriétaires(suite)/Owners(continued):**INTERNATIONAL BUSINESS MACHINES CORPORATION, US

(74) **Agent:** CHAN, BILL W.K.

(57) **Abrégé(suite)/Abstract(continued):**

system and method performs combining the confidence score with a degree of match between a LAT and an AnswerType of the candidate answer to provide improved overall score for the candidate answer.

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

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
29 March 2012 (29.03.2012)

(10) International Publication Number
WO 2012/040350 A1

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

(21) International Application Number:
PCT/US2011/052586

(22) International Filing Date:
21 September 2011 (21.09.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/386,029 24 September 2010 (24.09.2010) US

(71) Applicant (for all designated States except US): **INTERNATIONAL BUSINESS MACHINES CORPORATION** [US/US]; New Orchard Road, Armonk, New York 10504 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **FAN, James, J.** [US/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US). **FERRUCCI, David, A.** [US/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US). **GONDEK, David, C.** [US/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US). **KALYANPUR, Aditya, A.** [IN/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US). **LALLY, Adam, P.** [US/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US). **MURDOCK IV, James, W.** [US/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US). **ZADROZNY, Wlodek, W.** [US/US]; 19 Skyline Drive, Hawthorne, New York 10532 (US).

(74) Agents: **GROLZ, Edward, W.** et al.; Scully, Scott, Murphy & Presser, 400 Garden City Plaza, Suite 300, Garden City, New York 11530 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, 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, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, 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, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

[Continued on next page]

(54) Title: LEXICAL ANSWER TYPE CONFIDENCE ESTIMATION AND APPLICATION




Fig. 1

(57) Abstract: A system, method and computer program product for automatically estimating the confidence of a detected LAT to provide a more accurate overall score for an obtained candidate answer. A confidence "score" or value of each detected LAT is obtained, and the system and method performs combining the confidence score with a degree of match between a LAT and an AnswerType of the candidate answer to provide improved overall score for the candidate answer.

WO 2012/040350 A1

— *before the expiration of the time limit for amending the* *amendments (Rule 48.2(h))*
claims and to be republished in the event of receipt of

LEXICAL ANSWER TYPE CONFIDENCE ESTIMATION AND APPLICATION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present invention relates to and claims the benefit of the filing date of commonly-owned, co-pending United States Provisional Patent Application No. 61/386,029, filed September 24, 2010.

BACKGROUND

[0002] An introduction to the current issues and approaches of question answering (QA) can be found in the web-based reference http://en.wikipedia.org/wiki/Question_answering. Generally, QA is a type of information retrieval. Given a collection of documents (such as the World Wide Web or a local collection) the system should be able to retrieve answers to questions posed in natural language. QA is regarded as requiring more complex natural language processing (NLP) techniques than other types of information retrieval such as document retrieval, and it is sometimes regarded as the next step beyond search engines.

[0003] QA research attempts to deal with a wide range of question types including: fact, list, definition, How, Why, hypothetical, semantically-constrained, and cross-lingual questions. Search collections vary from small local document collections, to internal organization documents, to compiled newswire reports, to the World Wide Web.

[0004] Closed-domain QA deals with questions under a specific domain, for example medicine or automotive maintenance, and can be seen as an easier task because NLP systems can exploit domain-specific knowledge frequently formalized in ontologies. Open-domain QA deals with questions about nearly everything, and can only rely on general ontologies and world knowledge.

On the other hand, these systems usually have much more data available from which to extract the answer.

[0005] Alternatively, closed-domain QA might refer to a situation where only a limited type of questions are accepted, such as questions asking for descriptive rather than procedural information.

[0006] Access to information is currently dominated by two paradigms. First, a database query that answers questions about what is in a collection of structured records. Second, a search that delivers a collection of document links in response to a query against a collection of unstructured data, for example, text or html.

[0007] A major unsolved problem in such information query paradigms is the lack of a computer program capable of accurately answering factual questions based on information included in a collection of documents that can be either structured, unstructured, or both. Such factual questions can be either broad, such as "what are the risks of vitamin K deficiency?", or narrow, such as "when and where was Hillary Clinton's father born?"

[0008] It is a challenge to understand the query, to find appropriate documents that might contain the answer, and to extract the correct answer to be delivered to the user.

SUMMARY

[0009] In one aspect there is provided a computing infrastructure and methodology that conducts question and answering that addresses the problem of computing wrong LATs for a particular query.

[0010] In one aspect, a system, method and computer program product is provided that automatically estimates the confidence of a query detected LAT to provide a more accurate overall score for the candidate answer. A confidence "score" or value of each detected LAT is

obtained, and the system and method performs combining the confidence score with a degree of match between a LAT and a candidate answer to provide improved overall score for the candidate answer.

[0011] Thus, in one aspect, there is provided a system and method for automatically generating answers to questions. The method comprises: determining a lexical answer type (LAT) associated with an input query; computing a first score for the input query LAT, wherein the first score is indicative of a quality of the query LAT; obtaining a candidate answer to the input query from an information source; determining an answer type (AT) for the candidate answer; comparing the query LAT with the candidate answer AT; computing a second score representing a degree of match between the compared query LAT with the candidate answer AT; and combining the first score and the second score to provide a total score indicative of a quality of the candidate answer, wherein a processing device automatically performs one or more of the determining a query LAT, computing the first score and second score, obtaining candidate answers, determining a AT, comparing and combining

[0012] In one aspect, the computing a first score for the query LAT associated with the input query comprises: extracting a set of features describing the determined query LAT; and, applying a model to the feature set to produce a LAT confidence value representing an estimated accuracy of the determined query LAT, the first score comprising the confidence value.

[0013] In a further aspect, there is provided a system for automatically generating answers to questions comprising: a memory storage device; one or more processor devices, each in communication to the memory device and configured for performing a method comprising: determining a lexical answer type (LAT) associated with an input query; computing a first score for the query LAT, wherein the first score is indicative of a quality of the query LAT; obtaining a candidate answer to the input query from an information source; determining an answer type (AT) for the candidate answer; comparing the query LAT with the candidate answer AT; computing a second score representing a degree of match between the compared query LAT with

the candidate answer AT; and, combining the first score and the second score to provide a total score indicative of a quality of the candidate answer.

[0014] In yet a further aspect, there is provided a method for extracting features from a query comprising a text string. The method comprises: identifying a syntactic pattern rule associated with the query, the pattern rule identified from a set of pattern rules that define common lexical answer types (LATs), a first feature of the extracted features comprising an identified pattern rule; checking for prior instances of a detected lexical answer type (LAT) and computing a LAT word frequency based on the prior instances, a second feature of the extracted features comprising a computed frequency of a query word being a candidate LAT; obtaining a parse tree data structure associated with the query; identifying grammatical relations amongst words associated with the candidate LAT in the parse tree structure, a third feature of the extracted features comprising a part of speech of the candidate LAT, and determining whether the candidate LAT word co-references some other word in the query recognized as a LAT, a fourth feature of the extracted features comprising a co-reference information, wherein one or more programmed processor devices performs the identifying a syntactic pattern rule, checking for prior instances, obtaining the parse tree data structure, identifying grammatical relations, and determining LAT word co-references.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The objects, features and advantages of the invention are understood within the context of the Detailed Description, as set forth below. The Detailed Description is understood within the context of the accompanying drawings, which form a material part of this disclosure, wherein:

[0016] FIG. 1 shows a high level logical architecture 10 of a question/ answering method in which the present invention may be employed;

[0017] FIG. 2A is a flow diagram of a LAT confidence estimation process and software architecture for generating a LAT confidence estimation score;

[0018] FIG. 2B is a flow diagram depicting an application of a LAT confidence estimation;

[0019] FIG. 3 is a flow diagram illustrating feature(s) generation production performed at step 160 of the flow diagram of FIG. 2A;

[0020] FIG. 4 depicts an example portion of a decision tree model 190 that is traversed for computing a query LAT confidence score;

[0021] FIG. 5 shows a further high level logical architecture 10 of a question/ answering method in which the present invention may be employed;

[0022] FIG. 6 illustrates an example flow diagram for conducting questions and answers processing with deferred type evaluation;

[0023] FIG. 7 is a flow diagram illustrating TyCor score production performed at step 132 of the flow diagram of FIG. 6 in one embodiment.

[0024] FIG. 8 depicts an aspect of a UIMA framework implementation for providing one type of analysis engine for processing CAS data structures; and

[0025] FIG. 9 illustrates an exemplary hardware configuration to run method steps described in FIGS. 2A, 2B, 3-4, 6-7 in one embodiment.

DETAILED DESCRIPTION

[0026] Commonly-owned, co-pending U.S. Patent Application 12/126,642, titled "SYSTEM AND METHOD FOR PROVIDING QUESTION AND ANSWERS WITH DEFERRED TYPE EVALUATION" describes a QA system and method in which answers are automatically generated for questions. More particularly, it involves comparing the

lexical types determined from the question to the lexical types associated with each candidate answer that is derived from a search. The lexical answer type requested by the question is referred to herein as a "lexical answer type" or "LAT." The lexical types that are associated with each candidate answer is referred to herein as an answer type or "AT".

[0027] Generally, a question answering (QA) system returns a precise answer to a given question in natural language along with justifications by consulting a possibly heterogeneous collection of structured, semi-structured and unstructured information resources. FIG. 1 illustrates the major components that comprise a canonical question answering system 10 and their workflow.

[0028] The question analysis component 20 receives a natural language question 19 (e.g., "Who is the 42nd president of the United States?") and analyzes the question to produce, minimally, the semantic type of the expected answer (in this example, "president"), and optionally other analysis results for downstream processing. A search component 30a utilizing a core search engine formulates queries from the output 29 of question analysis and consults various resources such as the World Wide Web 41 or one or more knowledge resources, e.g., databases, knowledge bases 42, to retrieve documents, passages, web-pages, database tuples, etc., that are relevant to answering the question. The candidate answer generation component 30b then extracts from the search results 45 potential (candidate) answers to the question, which are then scored and ranked by the answer selection component 50 to produce a final ranked list of answers with associated confidence scores.

[0029] In a QA system with deferred type evaluation, after the extraction of candidate answers (i.e., candidate answer instances to queries obtained from a data corpus or the Internet), there is typically performed answer-typing, i.e., a determination of associations between searched candidate answers and their lexical types. The determined candidate answer type is compared against a computed lexical answer type ("LAT") of the query which is detected by question analysis, and an evaluation is made in determining the correctness of the answer.

[0030] A Type Coercion (TyCor) process is one in which an answer type (AT) of a candidate answer is “coerced” to the question LAT based on several techniques. The coercion process may involve mapping the candidate answer to an instance in a structured resource, obtaining its types and mapping them to the LAT. The results of the “coercion” are referred to as a TyCor score.

[0031] A LAT of the question/query is the type (i.e. the descriptor) of the referent of the entity that is a valid answer to the question/query. In practice, LAT is the descriptor of the answer detected by a natural language understanding module comprising a collection of patterns and/or a parser with a semantic interpreter.

[0032] Thus, question and answer systems may employ a Lexical Answer Type (LAT) detection processing block 200 in the query analysis module 20 of FIG. 1. The LAT further represents the question terms that identify the semantic type of the correct answer. Thus, a LAT may be detectable in a question through pattern LAT detection rules. These rules are implemented and are encoded or learned by machines automatically through association rule learning. A natural language understanding model may implement these rules.

[0033] However, no natural language processing module is perfect and erroneous LATs are detected. Wrong LATs can seriously reduce the performance of QA systems because wrong types of answers will be judged as more likely to be the correct answers.

[0034] Further details on a QA system 10 are discussed with respect to FIG. 5. Evidence gathering and answer scoring functions 50 are programmed processing components that receive (1) a “lexical answer type” (LAT), e.g., a text string describing some concept and (2) “candidate answer” input(s), e.g., a text string describing some entity. The programmed processing components generate an output including a judgment (e.g., a score) representing to what degree the entity is an instance of the concept. As described in FIGs. 6 and 7 below, this is used as part of answer scoring functions, for example, based on computed candidate answer types (ATs).

[0035] As described below with respect to FIG. 3, after processing an input query to determine a LAT and searching and obtaining one or more candidate answers, there is performed for each candidate answer received the steps of matching the candidate against instances in a database which results in generating an output data structure, including the matched instances 132a; looking (searching) for evidence that the candidate answer has the required LAT and retrieving LT(s) associated with those instances in the knowledge base (KB) 132b; and attempting to match LAT(s) with LT(s) (e.g., as determined by a matching function that using a parser, a semantic interpreter and/or a simple pattern matcher) and producing a score representing the degree of match 132c. More particularly, the candidate AT and query LAT(s) are represented as lexical strings. Production of the score, referred to as a "TyCor" (Type Coercion) score, is comprised of the three steps: (1) candidate answer to instance matching, (2) instance to type association extraction, and (3) LAT to type matching. The score reflects the degree to which the candidate may be "coerced" to the LAT, where higher scores indicate a better coercion.

[0036] In one embodiment, the present disclosure extends and complements the effectiveness of the QA system and method described in co-pending U.S. Patent Application 12/126,642 by providing additional functionality to provide a better quality score for the candidate answer. Specifically, a QA system such as shown in FIG. 1 performs an automatic estimation of the confidence of a detected query LAT to produce a LAT confidence score. The LAT confidence score is combined with the degree of match between a LAT and a answer type (AT) of the candidate answer to provide a better overall score for the candidate answer. In other words, a confidence score or value of each detected LAT is obtained and the system and method performs combining the LAT confidence score with the TyCor score to improve the overall score for the candidate answer. The attendant advantages include a QA system will be more tolerant of LAT detection errors and candidate answers are more likely to be correct when LAT confidence is taken into account.

[0037] With respect to computing a LAT confidence score, reference is had with respect to FIG. 2A, which depicts the process of LAT confidence estimation such as may be performed as part of, or in conjunction with the LAT detection processing block 200 in query analysis module 20,

shown in FIG. 1. In the processing, after computing a lexical answer types (LAT) for a query using the process as described in commonly-owned, co-pending U.S. Patent Application 12/126,642, the confidence “score” or confidence value of each LAT is computed.

[0038] FIG 2A particularly shows a computer-implemented process 150 for LAT confidence estimation including two components: a first feature generation component 160 that extracts a set of features 165 that describe a candidate LAT (e.g., a word in the question). Example features may include, but are not limited to: 1) the application of a pattern rule (“Rule ID”) from a set of rules that define common LATs, e.g., the rule with the id “thisWhatNoun” that describes syntactic patterns that match a noun modified by determiners such as “this” or “these”, as the noun “president” in “This president did what ...”; 2) a prior frequency of the word being a candidate LAT; 3) parse information to determine the grammatical relations associated with the LAT, e.g., how candidate LAT word relates to other parts of the parse tree; 4) Other grammatical information, such as part-of-speech of the LAT; and, 5) co-reference information that determines if the candidate LAT word co-references some other word recognized as a LAT, etc. The confidence estimation processing component 175 takes the set of LAT features, and a model 180 is applied to the set to produce a confidence score 199 associated with the LAT. The model 180 can be constructed manually or automatically from a set of training instances.

[0039] Thus for example, a pattern may be found in the question that leads to the LAT based on processing according to a Prolog[®] logic program in one embodiment. For example, appropriate weights for particular identified pattern(s) are determined by past training data, i.e., data used in previous processes for manually identifying LATs in the past. During LAT confidence estimation, the text surrounding the LAT is processed, e.g., using Prolog[®] logic program, to determine if any of the patterns fit. That is, a Prolog[®] logic program may be used in one embodiment to perform pattern rule matching. If a pattern fits, the ruleID associated with the identified pattern is provided and based on a prior confidence for that particular pattern, a weight is assigned to the feature. If none of the patterns match, there would be no weight associated with that feature.

[0040] An embodiment of the feature generation component 160 at the question analysis stage is now described with reference to FIG. 3. As shown, the query LAT is input to and processed in parallel in plural modules that each performs a specific feature extraction function. The plural modules include: a process 162 for identifying a pattern Rule ID from a set of rules that define common LATs; a process 164 for checking prior instances of a detected LAT word, e.g., from a collection of manually annotated data and computing a LAT word frequency based on an amount of prior instances; a parse tree checking process 166 to determine parts of speech of the candidate LAT (e.g., a word type such as a noun, verb, adjective, etc.); a process 168 for identifying grammatical relations associated with the LAT, e.g., how candidate LAT word relates to other parts of the parse tree; and a process 170 for identifying co-reference information that determines if the candidate LAT word co-references some other word recognized as a LAT (e.g., a question: “*This president, who was a general during the Revolutionary War...*”) provides two LATs (a president and a general) and if it is known that the word “president” is a LAT, then the word “general” is most likely a LAT as well as the candidate input LAT word “general” co-references the known LAT president. The output 172 of the processing blocks 162-170 provides a feature set having content including, but not limited to, one or more of: a Boolean (e.g., a grammatical relations, or a co-reference indicating whether the candidate LAT references another LAT or not); a numeric feature (e.g., word as LAT frequency); and a category (e.g., a part of speech such as a noun). The output feature set is then input to the confidence estimation module 175. The confidence estimation unit 175 receives the feature set, and the trained model 180 is applied to the feature set to produce an estimated LAT confidence score. In one embodiment, the feature set includes a list of numbers (e.g. 1, 0, 0.234, 1, 1, ...). Categorical features can be turned into a set of Boolean features (e.g. part-of-speech can be turned into is_noun, is_verb, is_adj, is_adv, is_prep).

[0041] In one embodiment, a feature set of Boolean, categorical values, and number values are input to the prior trained model. For the case of Boolean values or categorical values, for example, these may be represented in the feature set as a binary number such as 0 or 1. For example, a value of 0 or 1 may be associated with respective False or True feature set Boolean values.

[0042] With respect to the trained model 180, one embodiment is a machine learning model configured to combine the different features and weigh them to produce the final score for LAT confidence estimation. Any machine learning model may be used which is a mathematical model to determine how the different features of the feature set should be weighted. The mathematical model includes the weights for each type of feature, and given a particular input of feature set values, the values are input to the model and the model produces the score. The individual feature values are variables to the model equation (a function with different weights for each variable) and the application of the model a value is given to each of the variables to produce a number. The model can be constructed manually or automatically from a set of training instances.

[0043] In one embodiment, the trained model 180 is a logistic regression model a description of which is described in a reference entitled *Machine Learning*, Tom Mitchell, McGraw Hill, 1997.

In other embodiments, a support vector machine model may be used as described in a reference entitled "Support Vector Machines and other kernel-based learning methods" by John Shawe-Taylor and Nello Cristianini, Cambridge University Press, 2000. Further, a maximum entropy machine model may be found such as described in a reference entitled "A maximum entropy approach to natural language processing" by Adam L. Berger, et al. in Computational Linguistics Journal Volume 22 Issue 1, March 1996. Otherwise, a rule-based or decision-tree based model may be implemented.

[0044] An example logistic regression model that could be used shown as equation 1) as follows:

$$f(x_1, \dots, x_n) = \frac{1}{e^{-z} + 1} \quad 1)$$

where $x_1, \dots, x_n$ are feature values, and, in one embodiment,

$$z = c + \sum_{i=1}^n w_i x_i$$

where $a_1, \dots, a_n$ are coefficients and the constant c are assigned during the training of a model.

In one non-limiting example, a particular model may have coefficients assignments such as:

$a_1 = 0.484$, $x_1 = \text{thisWhatRule}$ fired

$a_2 = 2.8671$, $x_2 = \text{heShePronounRule}$ fired

$a_3 = 12.587$ $x_3 = \text{prior probability}$...

[0045] FIG. 4 shows a further example of a model in the form of a decision tree data structure 190 which includes a series of decision points (e.g., nodes 192, 194, 196 in the tree) and different decisions (branches in the tree). In the example shown, the decisions are whether the input fits a particular pattern or not. Each pattern/rule has an associated name, such as “thisWhatRule” which detects patterns such as “this X ...” in a question. If a pattern fits the input, then a particular confidence is returned as the LAT. Processing of decision tree data structure 190 includes automatically traversing the trained decision tree and computing a score at each decision step, e.g., steps 192, 194, 196, etc. that is used for computing the confidence score.

[0046] As mentioned, during QA processing, a determination is made as to whether the candidate answer is an instance of the query LAT. As the LAT detection may have errors, there is produced the estimated confidence score associated with the LAT to improve the results of the type-matching score (the TyCor), i.e., by combining the computed LAT confidence score with the raw TyCor score.

[0047] FIG. 2B describes a programmed process 250 for applying an estimated LAT confidence score 199 to produce a modified Tycor score. In the programmed LAT confidence application process 250 a single step is performed at a combining unit 275 that combines a LAT confidence estimate 199 and the raw Tycor score 260 to produce a final score 299. In one embodiment, the LAT confidence score may be a number from 0 to 1, and this may be multiplied by the raw TyCor score to compute the final TyCor score. The combination of LAT confidence score with raw Tycor score is implemented in a variety of ways, such as a confidence weighted average.

[0048] Given a query having one or more candidate LATs, the LAT confidence score is treated as a weight. For example, each LAT has a weight associated with it represented by the LAT confidence score. Given an answer candidate and associated raw TyCor score for each LAT, an overall score is estimated as the sum of each weight confidence multiplied by the type-matching (TyCor) score for each LAT. Thus, for the example query described herein above:

“This president, who was a general during the Revolutionary War...”

for the detected two LATs “president” and “general” a total confidence weighted average for the candidate answer “George Washington” is computed as a sum of a first total representing the weighted confidence score (e.g., 0.8) for “general” multiplied by its respective TyCor score (e.g., 1.0) combined with a second total representing the weighted confidence score for “president” (0.95) multiplied by its respective TyCor (e.g., 1.0). In general, given a query with one or more detected LATs, a total weighted confidence score is computed as follows:

Total confidence weighted average score =
 $(\text{confidence score}_{\text{LAT1}})(\text{TyCor}_{\text{LAT1}}) + (\text{confidence score}_{\text{LAT2}})(\text{TyCor}_{\text{LAT2}}) + \dots +$
 $(\text{confidence score}_{\text{LATn}})(\text{TyCor}_{\text{LATn}})$

[0049] This summation is computed to yield the confidence weighted average for “n” detected LATs of a single query.

[0050] FIG. 5 shows a QA system diagram described in U.S. Patent Application 12/126,642 depicting a high-level logical QA architecture 10. Generally, as shown in FIG. 5, the high level logical architecture 10 includes a Query Analysis module 20 implementing functions for receiving and analyzing a user query or question. The term “user” may refer to a person or persons interacting with the system, or refers to a computer system 22 generating a query by mechanical means, and where the term “user query” refers to such a mechanically generated query and context 19'. A candidate answer generation module 30 is provided to implement a search for candidate answers by traversing structured, semi structured and unstructured sources

contained in a Primary Sources module 11 and in an Answer Source Knowledge Base (KB) module 21 containing collections of relations and lists extracted from primary sources. All the sources of information can be locally stored or distributed over a network, including the Internet.

[0051] The Candidate Answer generation module 30 of architecture 10 generates a plurality of output data structures containing candidate answers based upon the analysis of retrieved data. In FIG. 5, an Evidence Gathering module 50 interfaces with the primary sources 11 and knowledge base 21 for concurrently analyzing the evidence based on passages having candidate answers, and scores each of candidate answers, in one embodiment, as parallel processing operations. In one embodiment, the architecture may be employed utilizing the Common Analysis System (CAS) candidate answer structures as is described in commonly-owned, issued U.S. Patent No. 7,139,752 .

[0052] As depicted in FIG. 5, the Evidence Gathering module 50 comprises a Candidate Answer Scoring module 40 for analyzing a retrieved passage and scoring each of candidate answers of a retrieved passage. The Answer Source Knowledge Base (KB) 21 may comprise one or more databases of structured or semi-structured sources (pre-computed or otherwise) comprising collections of relations (e.g., Typed Lists). In an example implementation, the Answer Source knowledge base may comprise a database stored in a memory storage system, e.g., a hard drive.

[0053] An Answer Ranking module 60 provides functionality for ranking candidate answers and determining a response 99 returned to a user via a user's computer display interface (not shown) or a computer system 22, where the response may be an answer, or an elaboration of a prior answer or request for clarification in response to a question -when a high quality answer to the question is not found. A machine learning implementation is further provided where the "answer ranking" module 60 includes a trained model component (not shown) produced using a machine learning techniques from prior data.

[0054] The processing depicted in FIG. 5, may be local, on a server, or server cluster, within an enterprise, or alternately, may be distributed with or integral with or otherwise operate in conjunction with a public or privately available search engine in order to enhance the question answer functionality in the manner as described. Thus, the method may be provided as a computer program product comprising instructions executable by a processing device, or as a service deploying the computer program product. The architecture employs a search engine (e.g., a document retrieval system) as a part of Candidate Answer Generation module 30 which may be dedicated to searching the Internet, a publicly available database, a web-site (e.g., IMDB.com) or, a privately available database. Databases can be stored in any storage system, non-volatile memory storage systems, e.g., a hard drive or flash memory, and can be distributed over the network or not.

[0055] As mentioned, the system and method of FIG. 5 makes use of the Common Analysis System (CAS), a subsystem of the Unstructured Information Management Architecture (UIMA) that handles data exchanges between the various UIMA components, such as analysis engines and unstructured information management applications. CAS supports data modeling via a type system independent of programming language, provides data access through a powerful indexing mechanism, and provides support for creating annotations on text data, such as described in "Design and implementation of the UIMA Common Analysis System" in 2004, 43:3 IBM Systems Journal. It should be noted that the CAS allows for multiple definitions of the linkage between a document and its annotations, as is useful for the analysis of images, video, or other non-textual modalities (as taught in U.S. Patent 7,139,752)

[0056] In one embodiment, the UIMA may be provided as middleware for the effective management and interchange of unstructured information over a wide array of information sources. The architecture generally includes a search engine, data storage, analysis engines containing pipelined document annotators and various adapters. The UIMA system, method and computer program may be used to generate answers to input queries. The method includes inputting a document and operating at least one text analysis engine that comprises a plurality of coupled annotators for tokenizing document data and for identifying and annotating a particular

type of semantic content. Thus it can be used to analyze a question and to extract entities as possible answers to a question from a collection of documents.

[0057] As further shown in greater detail in the architecture diagram of FIG. 5, the "Query Analysis" module 20 receives an input that comprises the query 19 entered, for example, by a user via their web-based browser device. An input query 19 may comprise a text string. The query analysis block 20 includes additionally a Lexical Answer Type (LAT) block 200 that implements functions and programming interfaces to provide additional constraints on the answer type (LAT). The computation in the block 20 comprises but is not limited to the Lexical Answer Type. In one embodiment, the LAT detection block 200 may include the LAT confidence estimation functionality for the detected LAT.

[0058] As mentioned above, a LAT of the question/query is the type (i.e. the descriptor) of the referent of the entity that is a valid answer to the question. In practice, LAT is the descriptor of the answer detected by a natural language understanding module comprising a collection of patterns and/or a parser with a semantic interpreter.

[0059] With reference to the Lexical Answer Type (LAT) block 200, in the query analysis module 20 of FIG. 5, the LAT represents the question terms that identify the semantic type of the correct answer. In one embodiment, as known, a LAT may be detected in a question through pattern LAT detection rules. These rules are implemented and can be encoded manually or learned by machine automatically through association rule learning. In this case, the natural language understanding model can be limited to implementation the rules.

[0060] FIG.6 is a flow diagram of a computer programmed method 100 for conducting questions and answers with deferred type evaluation in one embodiment as described in 12/126,642. Generally, in the method of "deferred type evaluation" depicted in FIG. 6, a first processing step 112 represents the step of receiving, at a processing device, an input query, and generating a data structure, e.g., a CAS structure, including a question string and context for input to a Lexical

Answer Type (LAT) processing unit block 200 (FIG. 5) where, as indicated at step 115, the Query is analyzed and lexical answer type (LAT) and LAT confidence score is computed.

[0061] As a result of processing in the LAT processing component, as run at step 115, there is generated an output data structure, e.g., a CAS structure, including the computed LAT and possibly additional terms from the original question.

[0062] As result of processing in the LAT block 200 then, as typified at step 120, there is generated an output data structure, e.g., a CAS structure, including the computed original query (terms, weights) (as described in co-pending United States Patent Application No. 12/152,441.

[0063] Returning to FIG. 6, at processing step 122, there is performed the step of searching for candidate answer documents, and returning the results.

[0064] As a result of processing in a candidate answer generation module, as typified at step 122, there is generated an output data structure, e.g., a CAS structure, including all of the documents found from the data corpus (e.g., primary sources and knowledge base).

[0065] In FIG. 6, step 128, there is depicted the step of analyzing each document for a candidate answer to produce a set of candidate answers which may be output as a CAS structure using LAT (the lexical answer type).

[0066] For the example questions discussed herein, as a result of processing in the candidate answer generation module 30, as typified at step 132, FIG. 6, those candidate answers that are found in the document is checked against the query LAT requirement and will be returned as answer(s) at step 133 (based on their scores).

[0067] In candidate answer to instance matching, the candidate is matched against an instance or instances within the knowledge resource, where the form the instance takes depends on the knowledge resource. With a structured knowledge base, instances may be entities, with an encyclopedic source such as Wikipedia instances may be entries in the encyclopedia, with lexical resources such as the WordNet[®] lexical database (A trademark of the Trustees of Princeton University) instances may be synset entries (sets of synonyms), and with unstructured document (or webpage) collections, instances may be any terms or phrases occurring within the text. If multiple instances are found, a rollup using an aggregation function is employed to combine the scores from all candidates. If no suitable instance is found, a score of 0 is returned.

[0068] Next, instance association information is extracted from the resource. This information associates each instance with a type or set of types. Depending on the resource, this may take different forms; in a knowledge base, this corresponds to particular relations of interest that relate instances to types, with an encyclopedic source, this could be lexical category information which assigns a lexical type to an entity, with lexical resources such as WordNet[®], this is a set of lexical relations, such as hyponymy, over synsets (e.g. "artist" is a "person"), and with unstructured document collections this could be co-occurrence or proximity to other terms and phrases representing type.

[0069] Then, each LAT is then attempted to match against each type. A lexical manifestation of the type is used. For example, with encyclopedias, this could be the string representing the category, with a lexical resource such as WordNet[®], this could be the set of strings contained within the synset. The matching is performed by using string matching or additional lexical resources such as Wordnet[®] to check for synonymy or hyponymy between the LAT and type. Special logic may be implemented for types of interest; for example person matcher logic may be activated which requires not a strict match, synonym, or hyponym relation, but rather that both LAT and type are hyponyms of the term "person". In this way, "he" and "painter", for example, would be given a positive score even though they are not strictly synonyms or hyponyms. Finally, the set of pairs of scores scoring the degree of match may be resolved to a single final score via an aggregation function.

[0070] Thus, an implementation set forth in steps 132a-132c of FIG. 7, for an example question, each candidate answer is automatically checked against the LAT requirement. This may be performed by the Candidate Answer Scoring block 40, shown in FIG. 5, as part of the evidence gathering module 50, and particularly, a Candidate Answer Type Analysis module 400 that produces a probability measure that Candidate Answer is of the correct type based, e.g., on a grammatical and semantic analysis of the document with which the Candidate Answer appears. As described, this processing entails using an automated scoring function that can be expressed as a weighted combination of different typing scores, and, in one example embodiment, it may be expressed as

$$\text{TyCorScore} = 0.2 * \text{TyCorWordNet} + 0.5 * \text{TyCorKB} + 0.4 * \text{TyCorDoc}$$

[0071] Other combinations of scores are possible, and the optimal scoring function can be learned as described in the co-pending United States Patent Application No.12/152,441, filed May 14, 2008.

[0072] The scoring function itself is a mathematical expression, that-- in one embodiment -- could be based on the logistic regression function (a composition of linear expressions with the exponential function), and may be applied to a much larger number of typing scores.

[0073] The output of the "Candidate Answer Scoring" module 40 is a CAS structure having a list of answers with their scores given by the processing modules in the answer scoring modules included in the Candidate Answer Scoring block 40 of the evidence gathering module 50. In one embodiment, these candidate answers are provided with the TyCor matching score combined with the weighted confidence score as described herein above.

[0074] Finally, returning to FIG. 6, at step 133, the top candidate answers (based on their TyCor scores) are returned and / or further processed for combination with respective computed weighted LAT confidence scores in the manner as described herein.

[0075] In one embodiment, the above-described modules of FIGs. 1 and 5 can be represented as functional components in UIMA is preferably embodied as a combination of hardware and software for developing applications that integrate search and analytics over a combination of structured and unstructured information. The software program that employs UIMA components to implement end-user capability is generally referred to as the application, the application program, or the software application.

[0076] The UIMA high-level architecture, one embodiment of which is illustrated in FIG. 5, defines the roles, interfaces and communications of large-grained components that cooperate to implement UIM applications. These include components capable of analyzing unstructured source artifacts, such as documents containing textual data and/or image data, integrating and accessing structured sources and storing, indexing and searching for artifacts based on discovered semantic content.

[0077] Although not shown, a non-limiting embodiment of the UIMA high-level architecture includes a Semantic Search Engine, a Document Store, at least one Text Analysis Engine (TAE), at least one Structured Knowledge Source Adapter, a Collection Processing Manager, at least one Collection Analysis Engine, all interfacing with Application logic. In one example embodiment, the UIMA operates to access both structured information and unstructured information to generate candidate answers and an answer in the manner as discussed herein. The unstructured information may be considered to be a collection of documents, and can be in the form of text, graphics, static and dynamic images, audio and various combinations thereof.

[0078] Aspects of the UIMA are further shown in FIG. 8, where there is illustrated a Analysis Engine (AE) 600 that can be a component part of the Text Analysis Engine (TAE). Included in the AE 600 is a Common Analysis System (CAS) 610, an annotator 620 and a controller 630. A

second embodiment of a TAE (not shown) includes an aggregate Analysis Engine composed of two or more component analysis engines as well as the CAS, and implements the same external interface as the AE 600.

Common Analysis System 610

[0079] The Common Analysis System (CAS) 610 is provided as the common facility that all Annotators 620 use for accessing and modifying analysis structures. Thus, the CAS 610 enables coordination between annotators 620 and facilitates annotator 620 reuse within different applications and different types of architectures (e.g. loosely vs. tightly coupled). The CAS 610 can be considered to constrain operation of the various annotators.

[0080] The CAS 610 principally provides for data modeling, data creation and data retrieval functions. Data modeling preferably defines a tree hierarchy of (data) types, as shown in the example Table 1 provided below. The types have attributes or properties referred to as features. In preferred embodiments, there are a small number of built-in (predefined) types, such as integer (int), floating point (float) and string; UIMA also includes the predefined data type “Annotation”. The data model is defined in the annotator descriptor, and shared with other annotators. In the Table 1, some □types?that are considered extended from prior art unstructured information management applications to accommodate question answering in the preferred embodiment of the invention include:

<u>TYPE</u>	(or feature)	<u>TYPE's PARENT</u>	(or feature type)
Query Record		Top	
Query		Query Record	
Query Context		Query Record	
Candidate Answer Record		Annotation	
Candidate Answer		Candidate Answer Record	
Feature: CandidateAnswerScore		Float	

QueryLexical Answer Type	Annotation
CandidateAnswer AT	Annotation
Feature: TyCorScore	Float

TABLE 1

[0081] In Table 1, for example, all of the question answering types (list in the left column) are new types and extend either another new type or an existing type (shown in the right column). For example, both Query and Query Context are kinds of Query Record, a new type; while Candidate Answer Record extends the UIMA type Annotation, but adds a new feature CandidateAnswerScore which is a Float. In addition, Table 1 describes the query LAT as having a UIMA Annotation type; CandidateAnswerAT is also an Annotation, but with an additional feature TyCorScore of type Float.

[0082] CAS 610 data structures may be referred to as “feature structures”. To create a feature structure, the type must be specified (see TABLE 1). Annotations (and - feature structures) are stored in indexes.

[0083] The CAS 610 may be considered to be a collection of methods (implemented as a class, for example, in Java or C++) that implements an expressive object-based data structure as an abstract data type. Preferably, the CAS 610 design is largely based on a TAE Feature-Property Structure, that provides user-defined objects, properties and values for flexibility, a static type hierarchy for efficiency, and methods to access the stored data through the use of one or more iterators.

[0084] The abstract data model implemented through the CAS 610 provides the UIMA 100 with, among other features: platform independence (i.e., the type system is defined declaratively, independently of a programming language); performance advantages (e.g., when coupling annotators 610 written in different programming languages through a common data model); flow

composition by input/output specifications for annotators 610 (that includes declarative specifications that allow type checking and error detection, as well as support for annotators (TAE) as services models); and support for third generation searching procedures through semantic indexing, search and retrieval (i.e. semantic types are declarative, not key-word based).

[0085] The CAS 610 provides the annotator 620 with a facility for efficiently building and searching an analysis structure. The analysis structure is a data structure that is mainly composed of meta-data descriptive of sub-sequences of the text of the original document. An exemplary type of meta-data in an analysis structure is the annotation. An annotation is an object, with its own properties, that is used to annotate a sequence of text. There are an arbitrary number of types of annotations. For example, annotations may label sequences of text in terms of their role in the document's structure (e.g., word, sentence, paragraph etc), or to describe them in terms of their grammatical role (e.g., noun, noun phrase, verb, adjective etc.). There is essentially no limit on the number of, or application of, annotations. Other examples include annotating segments of text to identify them as proper names, locations, military targets, times, events, equipment, conditions, temporal conditions, relations, biological relations, family relations or other items of significance or interest.

[0086] Typically an Annotator's 620 function is to analyze text, as well as an existing analysis structure, to discover new instances of the set of annotations that it is designed to recognize, and then to add these annotations to the analysis structure for input to further processing by other annotators 620.

[0087] In addition to the annotations, the CAS 610 of FIG. 7 may store the original document text, as well as related documents that may be produced by the annotators 620 (e.g., translations and/or summaries of the original document). Preferably, the CAS 610 includes extensions that facilitate the export of different aspects of the analysis structure (for example, a set of annotations) in an established format, such as XML.

[0088] More particularly, the CAS 610 is that portion of the TAE that defines and stores annotations of text. The CAS API is used both by the application and the annotators 620 to create and access annotations. The CAS API includes, preferably, at least three distinct interfaces. A Type system controls creation of new types and provides information about the relationship between types (inheritance) and types and features. One non-limiting example of type definitions is provided in TABLE 1. A Structure Access Interface handles the creation of new structures and the accessing and setting of values. A Structure Query Interface deals with the retrieval of existing structures.

[0089] The Type system provides a classification of entities known to the system, similar to a class hierarchy in object-oriented programming. Types correspond to classes, and features correspond to member variables. Preferably, the Type system interface provides the following functionality: add a new type by providing a name for the new type and specifying the place in the hierarchy where it should be attached; add a new feature by providing a name for the new feature and giving the type that the feature should be attached to, as well as the value type; and query existing types and features, and the relations among them, such as “which type(s) inherit from this type”.

[0090] Preferably, the Type system provides a small number of built-in types. As was mentioned above, the basic types are int, float and string. In a Java implementation, these correspond to the Java int, float and string types, respectively. Arrays of annotations and basic data types are also supported. The built-in types have special API support in the Structure Access Interface.

[0091] The Structure Access Interface permits the creation of new structures, as well as accessing and setting the values of existing structures. Preferably, this provides for creating a new structure of a given type; getting and setting the value of a feature on a given structure; and accessing methods for built-in types. Feature definitions are provided for domains, each feature having a range.

[0092] In an alternative environment, modules of FIGs. 1, 5 can be represented as functional components in GATE (General Architecture for Text Engineering) (see: <http://gate.ac.uk/releases/gate-2.0alpha2-build484/doc/userguide.html>). Gate employs components which are reusable software chunks with well-defined interfaces that are conceptually separate from GATE itself. All component sets are user-extensible and together are called CREOLE - a Collection of REusable Objects for Language Engineering. The GATE framework is a backplane into which plug CREOLE components. The user gives the system a list of URLs to search when it starts up, and components at those locations are loaded by the system. In one embodiment, only their configuration data is loaded to begin with; the actual classes are loaded when the user requests the instantiation of a resource.). *GATE components are one of three types of specialized Java Beans*: 1) Resource: The top-level interface, which describes all components. What all components share in common is that they can be loaded at runtime, and that the set of components is extendable by clients. They have Features, which are represented externally to the system as η -data?in a format such as RDF, plain XML, or Java properties. Resources may all be Java beans in one embodiment. 2) ProcessingResource: Is a resource that is runnable, may be invoked remotely (via RMI), and lives in class files. In order to load a PR (Processing Resource) the system knows where to find the class or jar files (which will also include the metadata); 3) LanguageResource: Is a resource that consists of data, accessed via a Java abstraction layer. They live in relational databases; and, VisualResource: Is a visual Java bean, component of GUIs, including of the main GATE gui. Like PRs these components live in .class or .jar files.

[0093] In describing the GATE processing model any resource whose primary characteristics are algorithmic, such as parsers, generators and so on, is modelled as a Processing Resource. A PR is a Resource that implements the Java Runnable interface. The GATE Visualisation Model implements resources whose task is to display and edit other resources are modelled as Visual Resources. The Corpus Model in GATE is a Java Set whose members are documents. Both Corpora and Documents are types of Language Resources(LR) with all LRs having a Feature Map (a Java Map) associated with them that stored attribute/valuc information about the resource. FeatureMaps are also used to associate arbitrary information with ranges of documents

(e.g. pieces of text) via an annotation model. Documents have a DocumentContent which is a text at present (future versions may add support for audiovisual content) and one or more AnnotationSets which are Java Sets.

[0094] As UIMA, GATE can be used as a basis for implementing natural language dialog systems and multimodal dialog systems having the disclosed question answering system as one of the main submodules. The above-referenced US Patent Nos. 6,829,603 and 6,983,252 and 7,136,909 enable one skilled in the art to build such an implementation.

[0095] FIG. 8 illustrates an exemplary hardware configuration of a computing system 401 in which the present system and method may be employed. The hardware configuration preferably has at least one processor or central processing unit (CPU) 411. The CPUs 411 are interconnected via a system bus 412 to a random access memory (RAM) 414, read-only memory (ROM) 416, input/output (I/O) adapter 418 (for connecting peripheral devices such as disk units 421 and tape drives 440 to the bus 412), user interface adapter 422 (for connecting a keyboard 424, mouse 426, speaker 428, microphone 432, and/or other user interface device to the bus 412), a communication adapter 434 for connecting the system 400 to a data processing network, the Internet, an Intranet, a local area network (LAN), etc., and a display adapter 436 for connecting the bus 412 to a display device 438 and/or printer 439 (e.g., a digital printer of the like).

[0096] As will be appreciated by one skilled in the art, aspects of the present invention may be embodied as a system, method or computer program product. Accordingly, aspects of the present invention may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present invention may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[0097] Any combination of one or more computer readable medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. A computer readable storage medium may be, for example, but not limited to, an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system, apparatus, or device, or any suitable combination of the foregoing. More specific examples (a non-exhaustive list) of the computer readable storage medium would include the following: an electrical connection having one or more wires, a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM or Flash memory), an optical fiber, a portable compact disc read-only memory (CD-ROM), an optical storage device, a magnetic storage device, or any suitable combination of the foregoing. In the context of this document, a computer readable storage medium may be any tangible medium that can contain, or store a program for use by or in connection with a system, apparatus, or device running an instruction.

[0098] A computer readable signal medium may include a propagated data signal with computer readable program code embodied therein, for example, in baseband or as part of a carrier wave. Such a propagated signal may take any of a variety of forms, including, but not limited to, electro-magnetic, optical, or any suitable combination thereof. A computer readable signal medium may be any computer readable medium that is not a computer readable storage medium and that can communicate, propagate, or transport a program for use by or in connection with a system, apparatus, or device running an instruction.

[0099] Program code embodied on a computer readable medium may be transmitted using any appropriate medium, including but not limited to wireless, wireline, optical fiber cable, RF, etc., or any suitable combination of the foregoing.

[0100] Computer program code for carrying out operations for aspects of the present invention may be written in any combination of one or more programming languages, including an object oriented programming language such as Java, Smalltalk, C++ or the like and conventional

procedural programming languages, such as the "C" programming language or similar programming languages. The program code may run entirely on the user's computer, partly on the user's computer, as a stand-alone software package, partly on the user's computer and partly on a remote computer or entirely on the remote computer or server. In the latter scenario, the remote computer may be connected to the user's computer through any type of network, including a local area network (LAN) or a wide area network (WAN), or the connection may be made to an external computer (for example, through the Internet using an Internet Service Provider).

[0101] Aspects of the present invention are described below with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. These computer program instructions may be provided to a processor of a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which run via the processor of the computer or other programmable data processing apparatus, create means for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks. These computer program instructions may also be stored in a computer readable medium that can direct a computer, other programmable data processing apparatus, or other devices to function in a particular manner, such that the instructions stored in the computer readable medium produce an article of manufacture including instructions which implement the function/act specified in the flowchart and/or block diagram block or blocks.

[0102] The computer program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other devices to cause a series of operational steps to be performed on the computer, other programmable apparatus or other devices to produce a computer implemented process such that the instructions which run on the computer or other

programmable apparatus provide processes for implementing the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0103] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of code, which comprises one or more operable instructions for implementing the specified logical function(s).

[0104] It should also be noted that, in some alternative implementations, the functions noted in the block may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be run substantially concurrently, or the blocks may sometimes be run in the reverse order, depending upon the functionality involved. It will also be noted that each block of the block diagrams and/or flowchart illustration, and combinations of blocks in the block diagrams and/or flowchart illustration, can be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0105] The embodiments described above are illustrative examples and it should not be construed that the present invention is limited to these particular embodiments. Thus, various changes and modifications may be effected by one skilled in the art without departing from the spirit or scope of the invention as defined in the appended claims.

CLAIMS

What is claimed is:

1. A computer-implemented method for automatically generating answers to questions comprising the steps of:
 - determining a lexical answer type (LAT) associated with an input query;
 - computing a first score for said input query LAT, wherein said first score is indicative of a quality of said query LAT;
 - obtaining a candidate answer to said input query from an information source;
 - determining an answer type (AT) for said candidate answer;
 - comparing said query LAT with said candidate answer AT;
 - computing a second score representing a degree of match between said compared query LAT with said candidate answer AT; and
 - combining said first score and said second score to provide a total score indicative of a quality of said candidate answer.
2. The computer-implemented method of Claim 1, wherein said computing a first score for said query LAT associated with said input query comprises:
 - extracting a set of features describing said determined query LAT; and,
 - applying a model to the feature set to produce a LAT confidence value representing an estimated accuracy of the determined query LAT, said first score comprising said confidence value.
3. The computer-implemented method of Claim 2, wherein said applied model is one or more of: a logistic regression model or a decision tree model constructed manually or automatically from prior sets of training instances.
4. The computer-implemented method of Claim 2, wherein a feature of said extracted feature set comprises a pattern rule, and wherein said extracting a set of features comprises:

identifying a syntactic pattern rule associated with said query, said pattern rule identified from a set of pattern rules that define common LATs.

5. The computer-implemented method of Claim 2, wherein a feature of said extracted feature set comprises a frequency of a query word being a candidate LAT, and wherein said extracting a set of features further comprises:

checking prior instances of a detected LAT word; and
computing a LAT word frequency based on said prior instances.

6. The computer-implemented method of Claim 2, wherein a feature of said extracted feature set comprises a part of speech of the candidate LAT, and said extracting a set of features further comprises:

obtaining a parse tree data structure associated with said query; and
traversing said parse tree to identify grammatical relations amongst words associated with said candidate LAT.

7. The computer-implemented method of Claim 2, wherein a feature of said extracted feature set comprises a co-reference information, and wherein said extracting a set of features further comprises: determining whether the candidate LAT word co-references some other word in said query recognized as a LAT.

8. The computer-implemented method of Claim 2, wherein one or more features of said extracted feature set are extracted in parallel.

9. The computer-implemented method of Claim 2, wherein said computing a second score representing a degree of match between said compared query LAT with said candidate answer AT includes:

matching said candidate answer against instances in a corpus or a knowledge base;
retrieving a type associated with an instance in the corpus or the knowledge base; and,

matching said query LAT with said retrieved type to produce a type coercion (TyCor) value, said second score comprising said TyCor value.

10. The computer-implemented method of Claim 9, wherein a query has up to n candidate LATs, said combining to provide said total score is a summation according to:

$$(\text{confidence score}_{\text{LAT1}})(\text{TyCor}_{\text{LAT1}}) + (\text{confidence score}_{\text{LAT2}})(\text{TyCor}_{\text{LAT2}}) + \dots + (\text{confidence score}_{\text{LATn}})(\text{TyCor}_{\text{LATn}})$$

wherein confidence score_{LAT1}, confidence score_{LAT2}, and confidence score_{LATn} are produced confidence values representing estimated accuracies of each determined LAT of said query, and TyCor_{LAT1}, TyCor_{LAT2}, and TyCor_{LATn} are the produced type coercion (TyCor) values of a candidate answer with each LAT.

11. A system for automatically generating answers to questions comprising:
a memory storage device;

one or more processor devices, each in communication to said memory device and configured for performing a method comprising:

determining a lexical answer type (LAT) associated with an input query;

computing a first score for said query LAT, wherein said first score is indicative of a quality of said query LAT;

obtaining a candidate answer to said input query from an information source;

determining an answer type (AT) for said candidate answer;

comparing said query LAT with said candidate answer AT;

computing a second score representing a degree of match between said compared query LAT with said candidate answer AT; and,
combining said first score and said second score to provide a total score indicative of a quality of said candidate answer.

12. The system of Claim 11, wherein said one or more processor devices for computing a first score for said query LAT is configured to carry out said computing step by:
- extracting a set of features describing said determined query LAT; and,
 - applying a model to the feature set to produce a LAT confidence value representing an estimated accuracy of the determined query LAT, said first score comprising said confidence value.
13. The system of Claim 12, wherein said one or more processor devices are configured to apply one or more of: a logistic regression model or a decision tree model constructed manually or automatically from prior sets of training instances.
14. The system of Claim 12, wherein a feature of said extracted feature set comprises a pattern rule, said one or more processor devices for extracting a set of features further performs: identifying a syntactic pattern rule associated with said query, said pattern rule identified from a set of pattern rules that define common LATs.
15. The system of Claim 12, wherein a feature of said extracted feature set comprises a frequency of a query word being a candidate LAT, said one or more processor devices for extracting a set of features further performs:
- checking prior instances of a detected LAT word; and
 - computing a LAT word frequency based on said prior instances.
16. The system of Claim 12, wherein a feature of said extracted feature set comprises a part of speech of the candidate LAT, said one or more processor devices for extracting a set of features further performs:
- obtaining a parse tree data structure associated with said query; and
 - traversing said parse tree to identify grammatical relations amongst words associated with said candidate LAT.

17. The system of Claim 12, wherein a feature of said extracted feature set comprises a co-reference information, said one or more processor devices for extracting a set of features further performs:

determining whether the candidate LAT word co-references some other word in said query recognized as a LAT.

18. The system of Claim 12, wherein said one or more processor devices for computing a second score is configured to carry out said computing step by:

matching said candidate answer against instances in a corpus or a knowledge base;
retrieving a type associated with an instance in the corpus or the knowledge base; and,
matching said query LAT with said retrieved type to produce a type coercion (TyCor) value, said second score comprising said TyCor value.

19. The system of Claim 18, wherein a query has n candidate LATs, said one or more processor devices for combining to provide said total score performs a summation according to:

$$(\text{confidence score}_{\text{LAT1}})(\text{TyCor}_{\text{LAT1}}) + (\text{confidence score}_{\text{LAT2}})(\text{TyCor}_{\text{LAT2}}) + \dots + (\text{confidence score}_{\text{LATn}})(\text{TyCor}_{\text{LATn}})$$

wherein confidence score_{LAT1}, confidence score_{LAT2}, and confidence score_{LATn} are produced confidence values representing estimated accuracies of each determined query LAT, and TyCor_{LAT1}, TyCor_{LAT2} and TyCor_{LATn} are the produced type coercion (TyCor) values of a candidate answer with each LAT.

20. A computer program product for automatically generating answers to questions, the computer program product comprising a non-transitory computer-readable medium readable by a processing circuit and storing instructions run by the processing circuit for performing a method, the method comprising:

determining a lexical answer type (LAT) associated with an input query;

computing a first score for said query LAT, wherein said first score is indicative of a quality of said query LAT;
obtaining a candidate answer to said input query from an information source;
determining an answer type (AT) for said candidate answer;
comparing said query LAT with said candidate answer AT; and,
computing a second score representing a degree of match between the compared query LAT with said candidate answer AT; and,
combining said first score and said second score to provide a total score indicative of a quality of said candidate answer.

21. The computer program product for automatically generating answers to questions of Claim 20, wherein said computing a first score for each LAT associated with said input query comprises:

extracting a set of features describing said query LAT; and,
applying a model to the feature set to produce a LAT confidence value representing an estimated accuracy of the determined query LAT, said first score comprising said LAT confidence value.

22. The computer program product for automatically generating answers to questions of Claim 21, wherein:

a first feature of said extracted feature set comprises a pattern rule, said extracting a set of features comprises identifying a syntactic pattern rule associated with said query, and said pattern rule is identified from a set of pattern rules that define common LATs; and wherein,

a second feature of said extracted feature set comprises a frequency of a query word being a candidate LAT, said extracting a set of features further comprises checking prior instances of a detected LAT word and computing a LAT word frequency based on said prior instances; and wherein,

a third feature of said extracted feature set comprises a part of speech of the candidate LAT, and said extracting a set of features further comprises obtaining a parse tree data structure

associated with said query, and traversing said parse tree to identify grammatical relations amongst words associated with said candidate LAT; and, wherein,

a fourth feature of said extracted feature set comprises a co-reference information, and said extracting a set of features further comprises determining whether the candidate LAT word co-references some other word in said query recognized as a LAT.





Fig. 1






FIG. 3




FIG. 4






Fig. 5

**Fig. 6**

**Fig. 7**

600 →




Fig. 8

FIG. 9

