



US009146917B2

(12) **United States Patent**
Barborak et al.

(10) **Patent No.:** **US 9,146,917 B2**
(45) **Date of Patent:** **Sep. 29, 2015**

(54) **VALIDATING THAT A USER IS HUMAN**

(75) Inventors: **Michael A. Barborak**, New York, NY (US); **David A. Ferrucci**, Yorktown Heights, NY (US); **James W. Murdock, IV**, Millwood, NY (US); **Wlodek W. Zadrozny**, Tarrytown, NY (US)

(73) Assignee: **International Business Machines Corporation**, Armonk, NY (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 27 days.

(21) Appl. No.: **13/547,436**

(22) Filed: **Jul. 12, 2012**

(65) **Prior Publication Data**

US 2013/0019285 A1 Jan. 17, 2013

Related U.S. Application Data

(60) Provisional application No. 61/508,287, filed on Jul. 15, 2011.

(51) **Int. Cl.**

G06F 21/31 (2013.01)
G06F 17/27 (2006.01)

(52) **U.S. Cl.**

CPC **G06F 17/27** (2013.01); **G06F 21/31** (2013.01); **G06F 2221/2103** (2013.01); **G06F 2221/2133** (2013.01)

(58) **Field of Classification Search**

CPC G06F 21/17
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

7,512,978 B1 3/2009 Screen et al.
7,680,891 B1 3/2010 Pongsajapan
7,841,940 B2* 11/2010 Bronstein 463/29

8,380,503 B2	2/2013	Gross
2005/0086045 A1	4/2005	Murata
2006/0286530 A1	12/2006	Forrest et al.
2008/0049939 A1	2/2008	Canetti et al.
2009/0287678 A1	11/2009	Brown et al.
2009/0292687 A1	11/2009	Fan et al.
2010/0205006 A1	8/2010	Bergh
2011/0119212 A1*	5/2011	De Bruin et al. 706/12
2011/0125734 A1*	5/2011	Duboue et al. 707/723
2012/0054834 A1	3/2012	King

OTHER PUBLICATIONS

Herdagdelen et al., "The Concept Game: Better Commonsense Knowledge Extraction by Combining Text Mining and a Game with a Purpose", Commonsense Knowledge Papers from the AAAI Fall Symposium, pp. 52-57, 2010.
Ethel Ong, "A Commonsense Knowledge Base for Generating Children's Stories", Commonsense Knowledge Papers from the AAAI Fall Symposium, pp. 82-87, 2010.
Chang et al., "Human Computation Games for Commonsense Data Verification", Commonsense Knowledge Papers from the AAAI Fall Symposium, pp. 19-20, 2010.
Bosch et al., "A Kid's Open Mind Common Sense", Commonsense Knowledge Papers from the AAAI Fall Symposium, pp. 114-119, 2010.

(Continued)

Primary Examiner — Samson Lemma

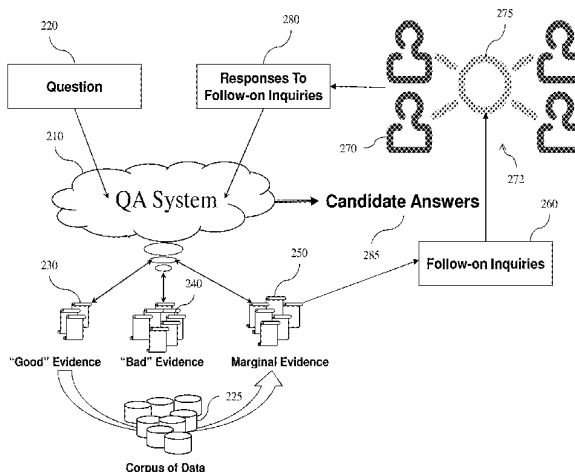
Assistant Examiner — Simon Kanaan

(74) *Attorney, Agent, or Firm* — Gibb & Riley, LLC

(57) **ABSTRACT**

A method of validating that a user is human. A user question is generated using a computerized device. The user question is output to a user. A user response to the user question is received from the user. The user response is validated as having been provided by a human.

24 Claims, 16 Drawing Sheets



(56)

References Cited

OTHER PUBLICATIONS

Havasi et al., "An Evaluation of the Brandeis Semantic Ontology", Computer Science Dept. Brandeis University, pp. 1-8, 2012.
Havasi et al., "Digital Intuition: Applying Common Sense Using Dimensionality Reduction", IEEE Intelligent Systems, 2009, pp. 1-12.
Lin et al., "Commonsense from the Web: Relation Properties", Commonsense Knowledge Papers from the AAAI Fall Symposium, pp. 70-75, 2010.
James Fan et al., "PRISMATIC: Inducing Knowledge from a Large Scale Lexicalized Relation Resource", Proceedings of the NAACL

HLT 2010 First International Workshop on Formalisms and Methodology for Learning by Reading, pp. 122-127, Los Angeles, California, Jun. 2010.
Office Action Communication, U.S. Appl. No. 13/564,023, Dated Mar. 18, 2014, pp. 1-32.
IBM Confidential Appendix P: List of IBM Patents or Patent Applications Treated As Related, I Page, Apr. 10, 2014.
Notice of Allowance Communication, U.S. Appl. No. 13/547,492, Dated Jan. 30, 2015, pp. 1-22.
Notice of Allowance Communication, U.S. Appl. No. 13/563,838, Dated Feb. 2, 2015, pp. 1-13.
Office Action Communication Dated Jun. 3, 2015, U.S. Appl. No. 13/570,680, pp. 1-23.

* cited by examiner

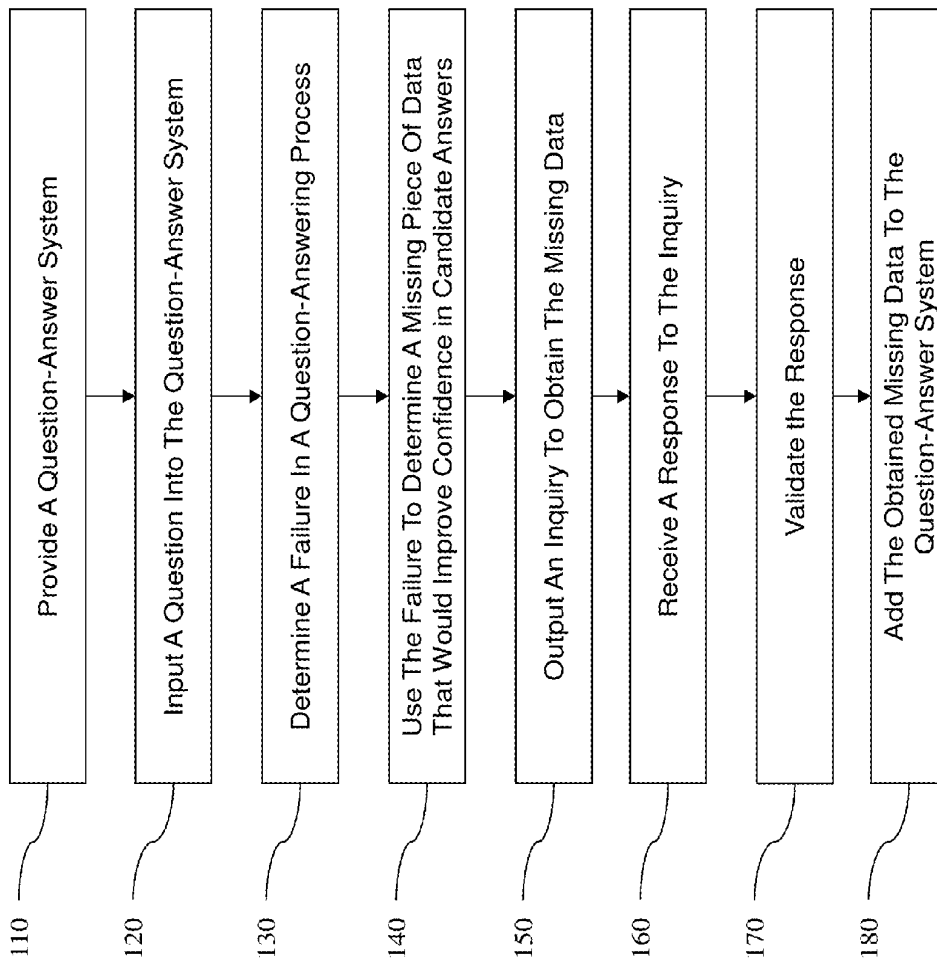


FIG 1

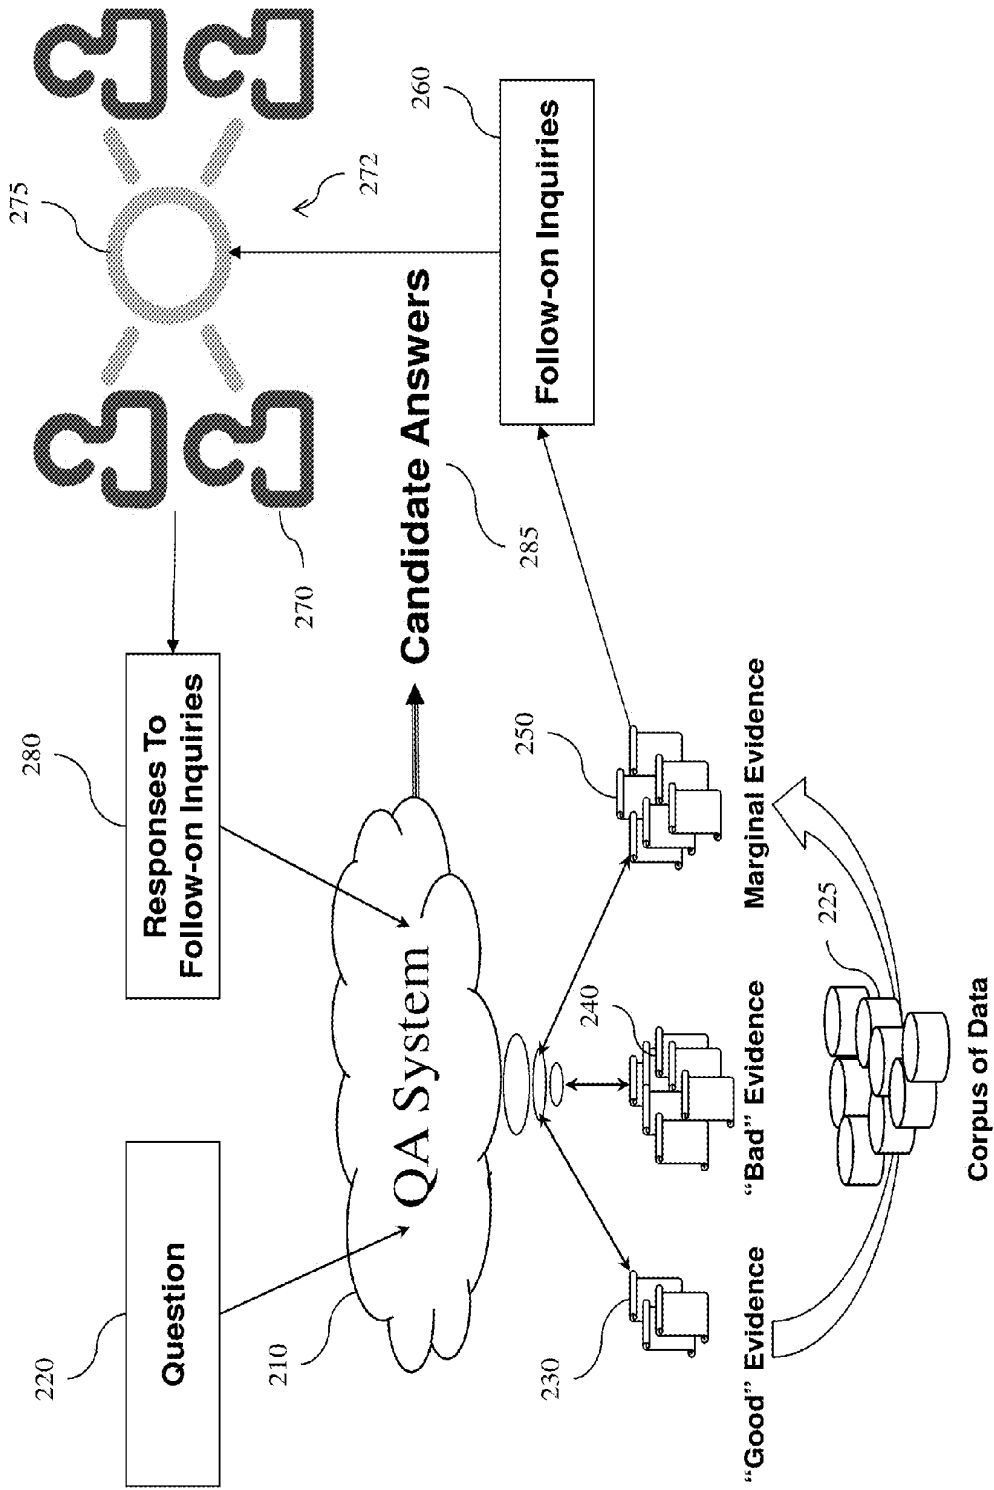


FIG. 2




FIG 3




FIG 4

510

Example 1: follow-on inquiries example can be Natural-language questions, e.g.:

- If I know that Eleazar stabbed an elephant from underneath, **can I conclude that Eleazar killed an elephant?**
- If I know that Eleazar stabbed an elephant from underneath, **would I suspect that Eleazar killed an elephant?**

515

Example 2: follow-on inquiries can be multiple-choice taxonomies, using text and/or diagrams:

- A) "kill" and "stab" are synonyms
- B) "kill" is a hyponym of "stab"
- C) "kill" is a hypernym of "stab"
- D) "kill" overlaps with "stab"
- E) None of the above

The diagrams show the following relationships:

- Diagram A: Two separate, non-overlapping ovals, one labeled 'kill' and one labeled 'stab', representing synonyms.
- Diagram B: A large oval labeled 'stab' containing a smaller oval labeled 'kill', representing 'kill' as a hyponym of 'stab'.
- Diagram C: A large oval labeled 'kill' containing a smaller oval labeled 'stab', representing 'kill' as a hypernym of 'stab'.
- Diagram D: Two overlapping ovals, one labeled 'kill' and one labeled 'stab', representing an overlap between the two terms.

FIG 5


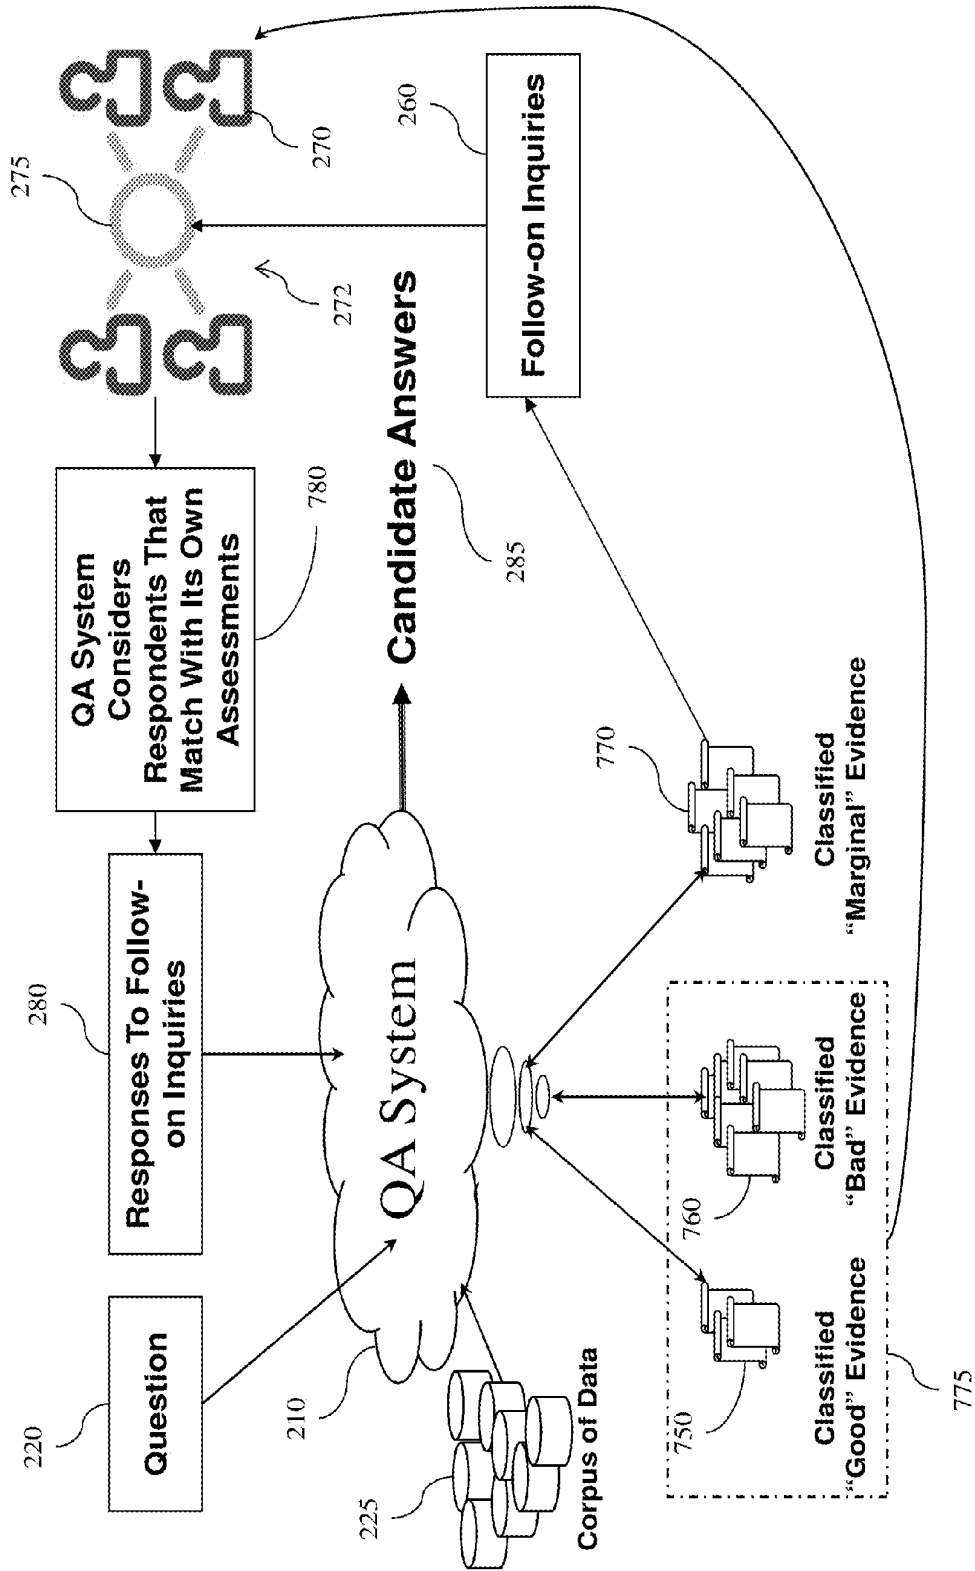


FIG 6






FIG 8

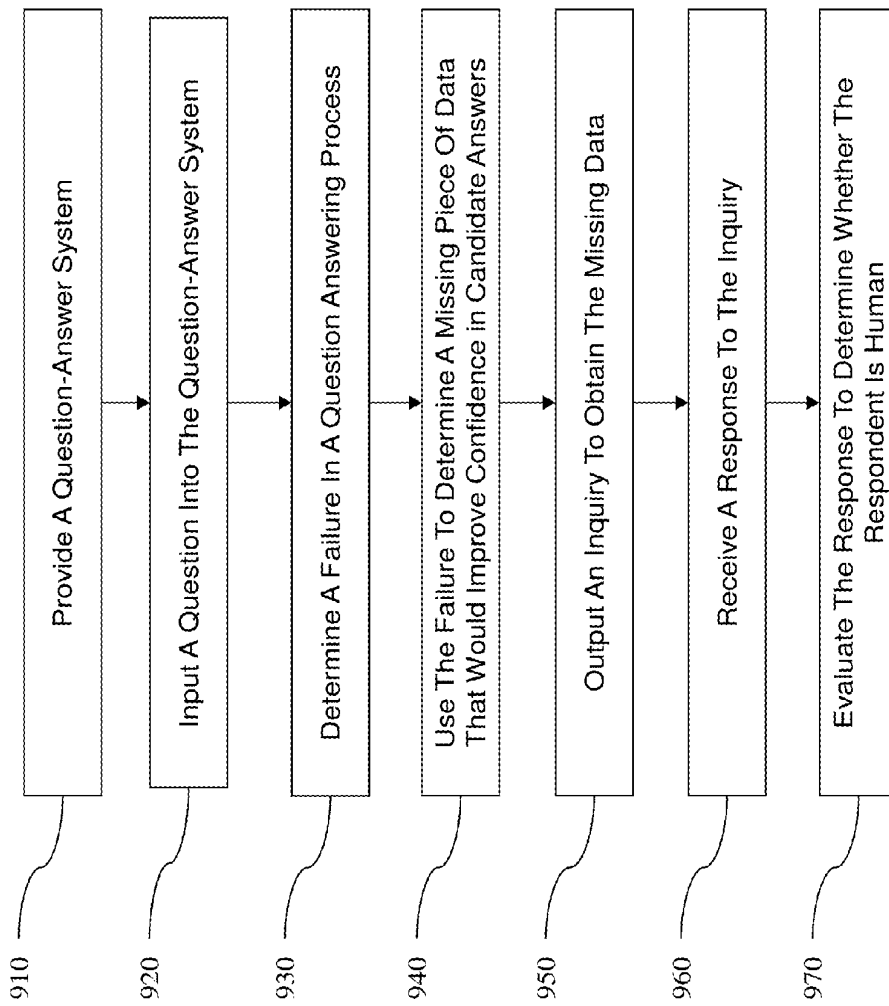


FIG 9




FIG 10




FIG 11

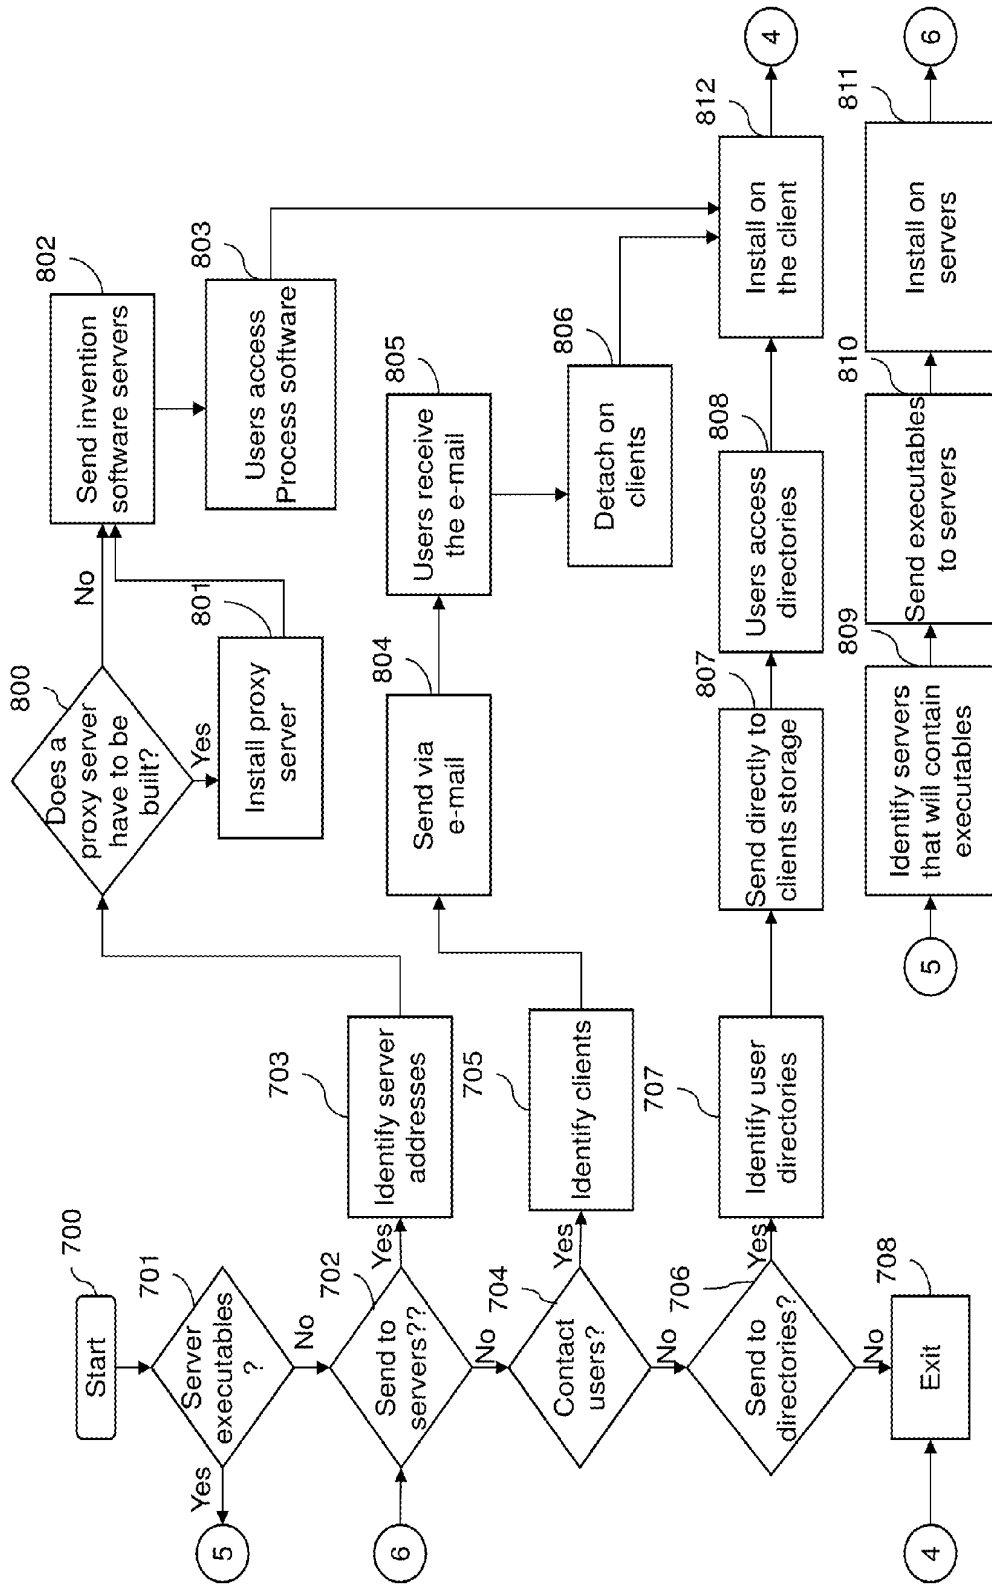


FIG 12

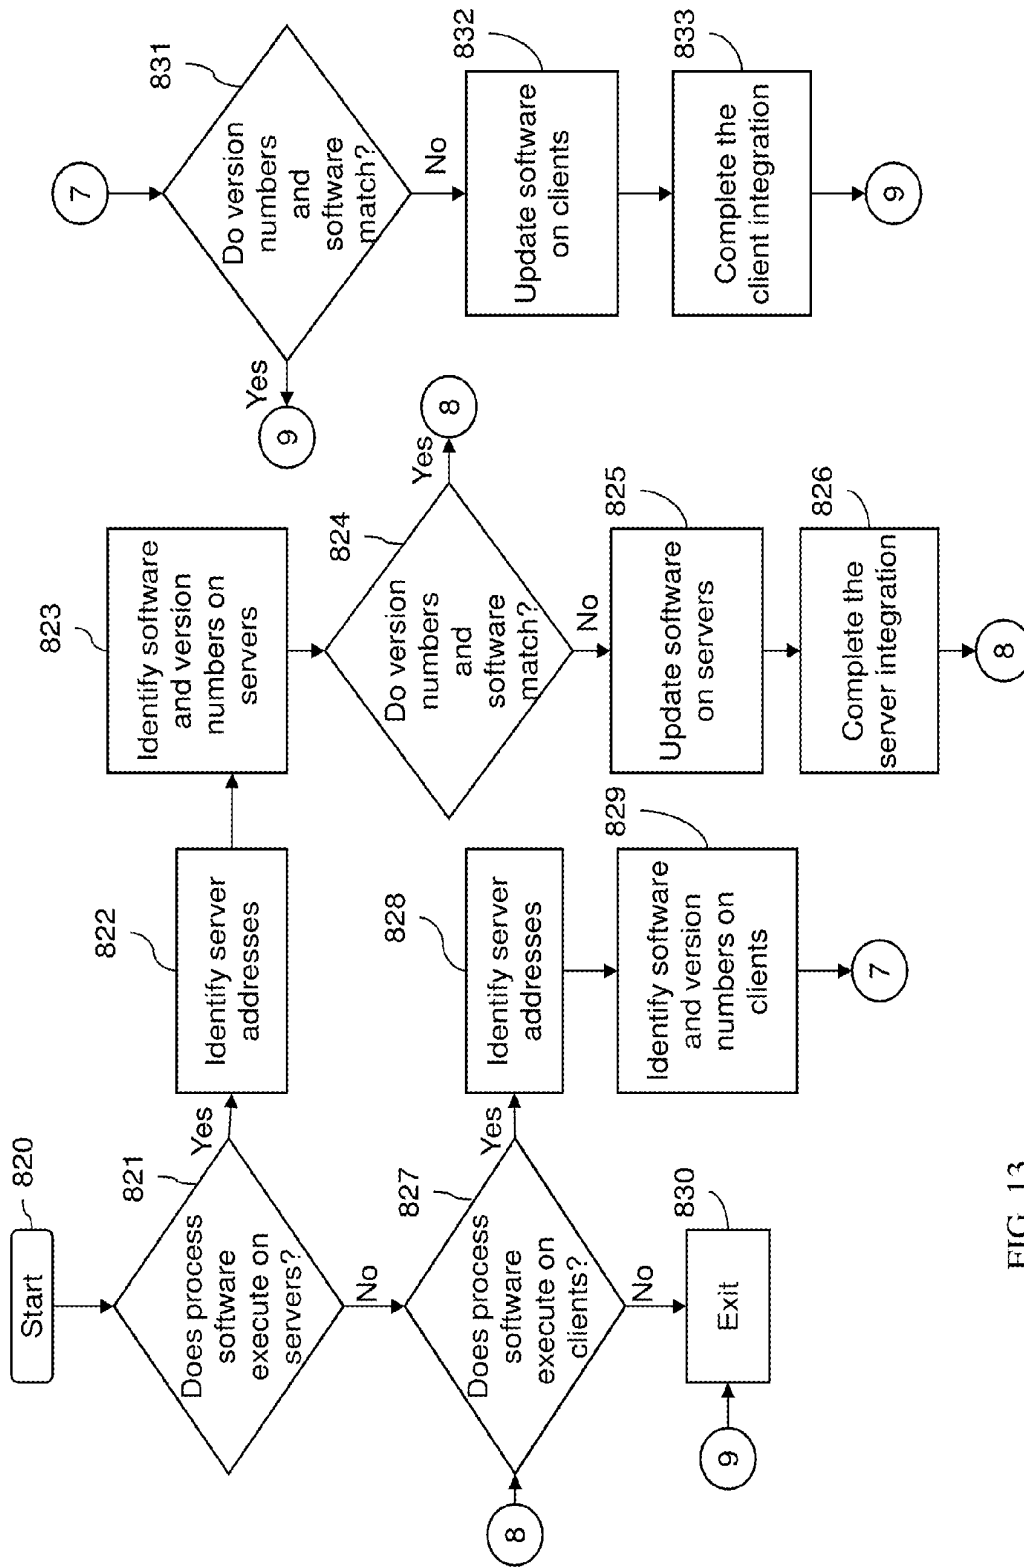


FIG 13

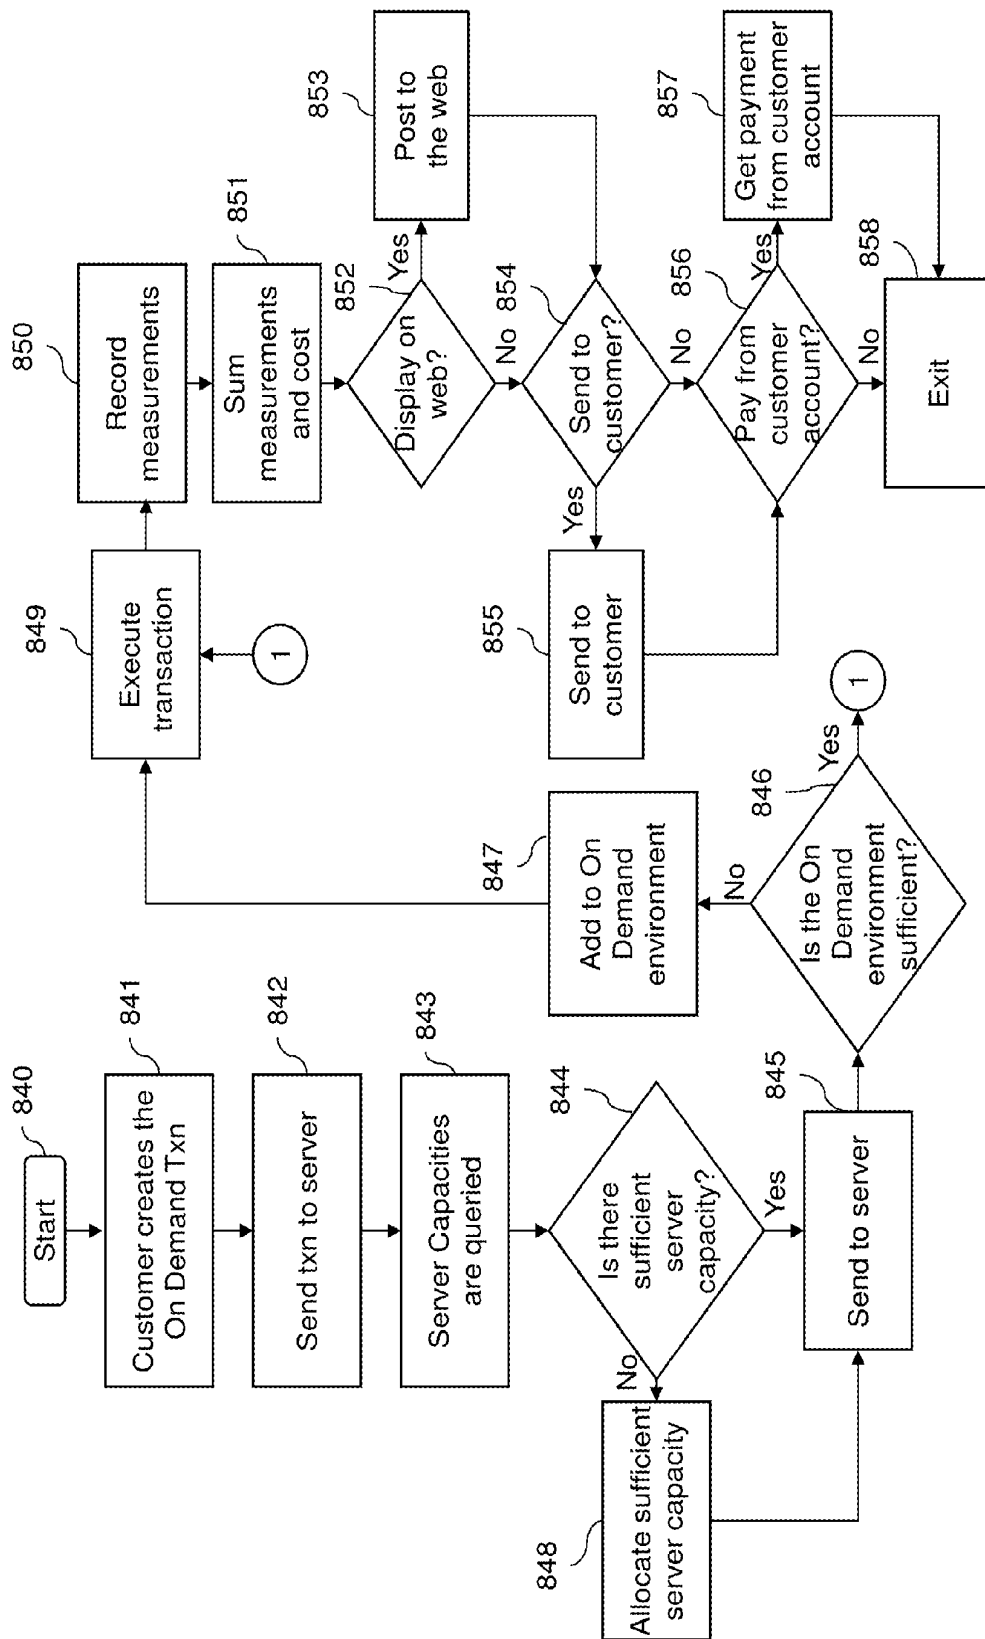


FIG 14

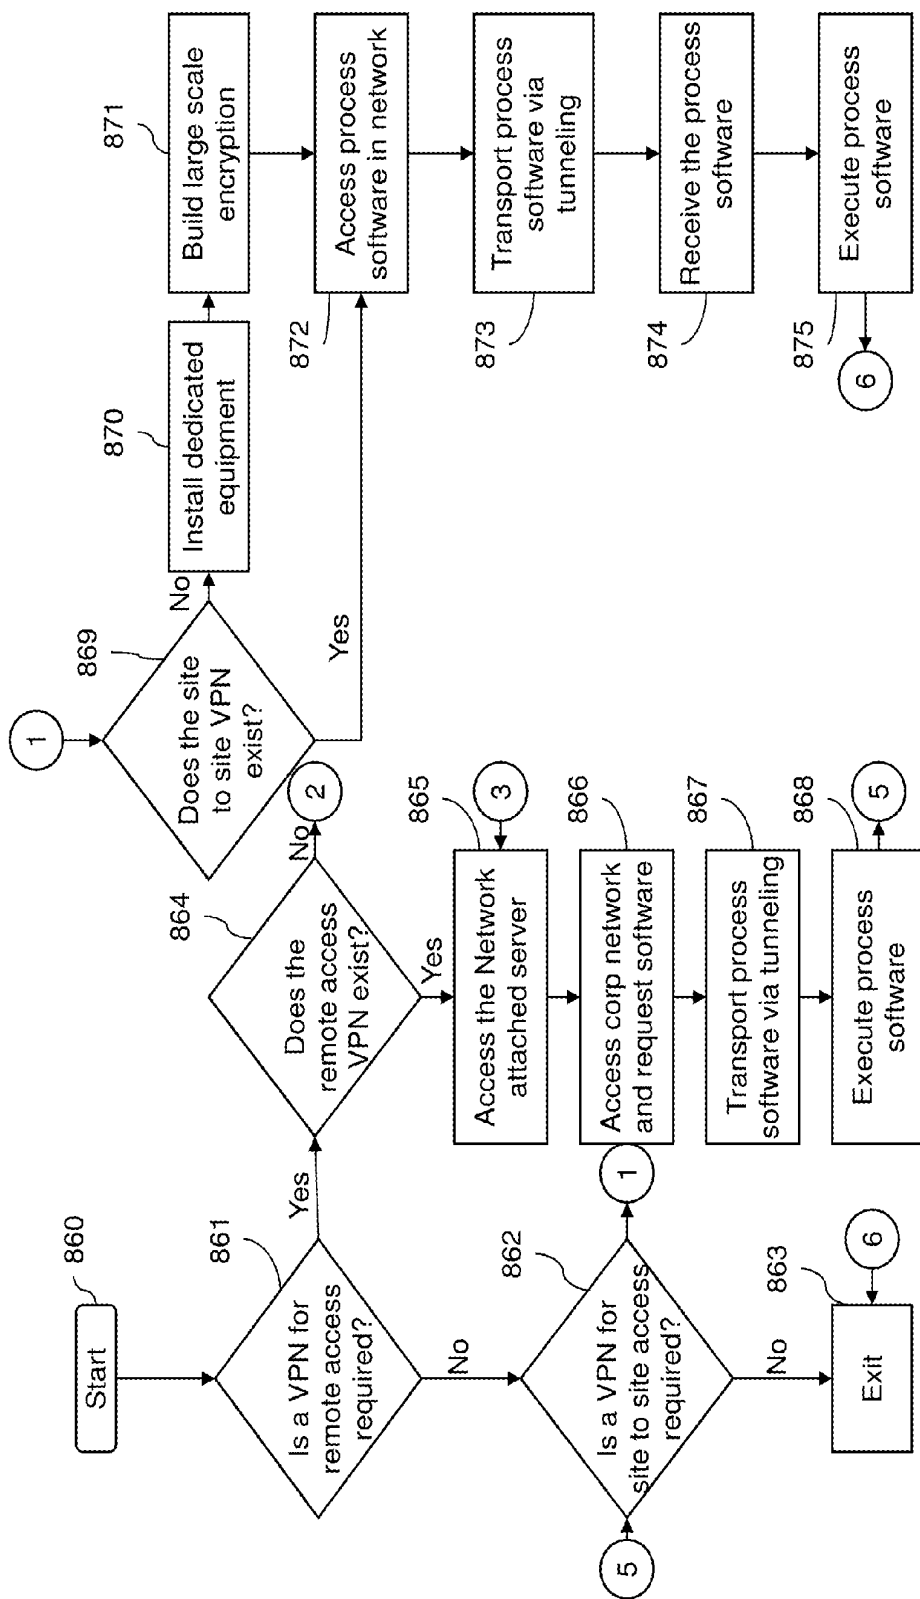


FIG 15




FIG 16

VALIDATING THAT A USER IS HUMAN**CROSS REFERENCES TO RELATED APPLICATIONS**

The present application claims the benefit under 35 U.S.C. §120 as a non-provisional of presently pending U.S. Patent Application Ser. No. 61/508,287 entitled "A METHOD FOR VALIDATING THAT A USER IS HUMAN", filed on Jul. 15, 2011, the entire teachings of which are incorporated herein by reference.

In addition, the present application is related to U.S. patent application Ser. No. 12/152,411 entitled, "System and Method for Providing Answers to Questions," filed on May 14, 2008, now U.S. Pat. No. 8,275,803, and to U.S. patent application Ser. No. 12/126,642 entitled, "System and Method for Providing Question and Answers with Deferred Type Evaluation," filed May 23, 2008, now U.S. Pat. No. 8,332,394, and to presently pending U.S. patent application Ser. No. 12/724,156, entitled, "Questions and Answers Generation," filed March 15, 2010, the entire teachings of each are incorporated herein by reference.

In addition, the present application is related to concurrently filed U.S. patent application Ser. No. 13/547,492, entitled, "A Method for Utilizing Failures in Question and Answer System Responses to Enhance the Accuracy of Question and Answer Systems".

BACKGROUND

The present disclosure relates to question answering (QA) systems, and more specifically, to validation that a user of a QA system is human.

Conventional technologies provide challenge systems for use with computer platforms to verify that a user of the computer platform is a human and not an automated response script. Typically, an audio or visual CAPTCHA device may be used for authentication or to prevent automated software from performing actions that are undesired. A CAPTCHA is a type of challenge-response test used in computing to ensure that the response is generated by a person instead of by an automated device. The process usually involves a computer asking a user to complete a simple test that the computer is able to recognize. The test is designed to be easy for a computer to generate, but difficult for a computer to solve. Therefore, if a correct solution is received, it can be presumed to have been entered by a human. A common type of CAPTCHA requires the user to type letters or digits from a distorted image that appears on the screen. Such tests are commonly used to prevent unwanted automated programs from accessing a computer platform.

SUMMARY

Embodiments herein disclose a mechanism for generating questions, presenting a question to a user, and using a response from the user to determine if the user is human.

Embodiments herein disclose a mechanism in which an automated question answering system generates secondary questions, asks those questions to a pool of users, and improves its ability to answer questions based on responses from those users.

According to an embodiment herein, a method of validating that a user is human is disclosed. A user question is generated using a computerized device. The user question is output to a user. A user response to the user question is

received from the user. The user response is validated as having been provided by a human.

According to another embodiment herein, a method of determining that a user of a question-answer system is human is disclosed. A question is received into a computerized question-answer system operating on a computerized device to provide an automatically generated answer to the question. A plurality of candidate answers to the question is generated by the computerized device. The sources of evidence used to generate the plurality of candidate answers are evaluated to identify marginal evidence. The marginal evidence contributes only partially to a candidate answer. The computerized device determines a confidence score for each of the plurality of candidate answers. The computerized device identifies information not provided by the marginal evidence that could further develop the confidence score. At least one follow-on inquiry is formulated based on the information. The follow-on inquiry is output to external sources separate from the question-answer system to obtain responses to the follow-on inquiry. At least one response to the inquiry is received from the external sources. Whether the at least one response to the at least one follow-on inquiry was provided by a human respondent is determined based on one of: a previous determination that a respondent is human, and syntactic or semantic characteristics of the at least one response.

According to another embodiment herein, a method is disclosed. A first question to be answered by a Question Answering (QA) system is provided to a processor. The processor creates a collection of candidate answers to the first question. The collection of candidate answers is created from a corpus of data. The processor generates supporting evidence and a confidence score for each candidate answer. An evidence analysis module produces a second question based on the supporting evidence. The processor presents the second question to one or more external sources separate from the QA system to obtain responses to the second question. The processor receives at least one response or knowledge item from the one or more external sources. The processor determines whether a respondent is human based on one of: a previous determination that a respondent is human, and syntactic or semantic characteristics of the at least one response or knowledge item provided to the second question.

According to another embodiment herein, another method is disclosed. A question is received into a Question Answering (QA) system. The QA system compares the question to a corpus of data. The QA system generates hypotheses about relationships between linguistic and semantic entities of the question and the corpus of data. The QA system generates a plurality of candidate answers to the question using the hypotheses. The QA system evaluates sources of evidence used to generate the plurality of candidate answers to identify marginal evidence. The marginal evidence contributes only partially to a candidate answer. The QA system determines a confidence score for each of the plurality of candidate answers. A missing piece of common-sense knowledge that would enable the QA system to further develop the confidence score is determined. A follow-on inquiry to obtain the missing piece of common-sense knowledge is formulated. The follow-on inquiry is output to an external expert community source separate from the QA system to obtain responses to the follow-on inquiry. Responses comprising the missing piece of common-sense knowledge are received from the external expert community source. A common-sense knowledge inquiry is generated. The common-sense knowledge inquiry is output to the external expert community source. A response to the common-sense knowledge inquiry is received from the external expert community source. Whether the

respondent is human is determined based on syntactic or semantic characteristics of the response to the common-sense knowledge inquiry.

According to another embodiment herein, a computerized device for determining that a user is human is disclosed. The computerized device comprises a question-answer system comprising software for performing a plurality of question answering processes. A receiver receives a question into the question-answer system. A processor that generates a plurality of candidate answers to the question is connected to the question-answer system. The processor evaluates sources of evidence used to generate the plurality of candidate answers to identify marginal evidence. The marginal evidence only partially contributes to a candidate answer. The processor determines a confidence score for each of the plurality of candidate answers. The processor identifies information not provided by the marginal evidence that could further develop the confidence score. The processor generates at least one follow-on inquiry based on the information. A network interface outputs the at least one follow-on inquiry to external sources separate from the question-answer system to obtain responses to the follow-on inquiry and receives from the external sources at least one response to the at least one follow-on inquiry. The processor evaluates whether the at least one response to the at least one follow-on inquiry was provided by a human respondent based on one of: a previous determination that a respondent is human and syntactic or semantic characteristics of the at least one response.

According to another embodiment herein, a computer system for verifying that a user is human is disclosed. The computer system comprises an automated question answering (QA) system comprising a corpus of data, a processor, a receiver, and a network interface. The processor has software for performing a plurality of question answering processes. The receiver receives a question into the automated QA system. The processor compares the question to the corpus of data and generates hypotheses about relationships between linguistic and semantic entities of the question and the corpus of data. The processor generates a plurality of candidate answers to the question using the hypotheses. The processor evaluates sources of evidence used to generate the plurality of candidate answers to identify marginal evidence. The marginal evidence only partially contributes to a candidate answer. The processor determines a confidence score for each of the plurality of candidate answers. The processor identifies a missing piece of common-sense knowledge; the missing piece of common-sense knowledge comprises information that would enable the QA system to further develop the confidence score. The processor automatically generates at least one secondary inquiry to obtain the missing piece of common-sense knowledge. The network interface outputs the at least one secondary inquiry to an external expert community source separate from the QA system and receives from the external expert community source at least one response to the secondary inquiry comprising the missing piece of common-sense knowledge. The processor generates a common-sense knowledge inquiry. The network interface outputs the common-sense knowledge inquiry to the external expert community source and receives from the external expert community source a response to the common-sense knowledge inquiry. The processor determines whether a respondent is human based on syntactic or semantic characteristics of the response to the common-sense knowledge inquiry.

According to another embodiment herein, a question answering (QA) system is disclosed. The QA system comprises a processor. An evidence analysis module is connected to the processor. A first interface is connected to the processor.

A second interface is connected to the processor. A corpus of data is connected to the evidence analysis module. The first interface provides a first question to be answered by the QA system. The processor creates a collection of candidate answers to the first question from the corpus of data. Each candidate answer has supporting evidence and a confidence score generated by the processor. The evidence analysis module produces a second question based on the supporting evidence. The processor presents the second question through the second interface to one or more external sources separate from the QA system to obtain responses to the second question. The processor receives at least one response or knowledge item from one or more external sources through the second interface. The processor determines whether a user is human based on one of: a previous determination that a respondent is human and syntactic or semantic characteristics of the at least one response or knowledge item obtained for the second question.

According to another embodiment herein, a non-transitory computer readable storage medium readable by a computerized device is disclosed. The computerized device comprises a question-answer system comprising question answering processes. The non-transitory computer readable storage medium stores instructions executable by the computerized device to perform a method. According to the method, a question is received into the question-answer system. A plurality of candidate answers to the question is generated. The sources of evidence used to generate the plurality of candidate answers are evaluated to identify marginal evidence. The marginal evidence contributes only partially to a candidate answer. A confidence score for each of the plurality of candidate answers is determined. The information not provided by the marginal evidence that could further develop the confidence score is identified. At least one follow-on inquiry is formulated based on the information. The follow-on inquiry is output to external sources separate from the question-answer system to obtain responses to the follow-on inquiry. At least one response to the follow-on inquiry is received from the external sources. Whether the at least one response to the at least one follow-on inquiry was provided by a human respondent is determined based on one of: a previous determination that a respondent is human and syntactic or semantic characteristics of the at least one response.

BRIEF DESCRIPTION OF THE DRAWINGS

The embodiments herein will be better understood from the following detailed description with reference to the drawings, which are not necessarily drawn to scale and in which:

FIG. 1 is a flow diagram of a method according to an embodiment herein;

FIG. 2 is an illustration of aspects of a system according to an embodiment herein;

FIG. 3 is an illustration of aspects of a system according to an embodiment herein;

FIG. 4 is an illustration of aspects of a system according to an embodiment herein;

FIG. 5 is an illustration of inquiry examples according to an embodiment herein;

FIG. 6 is an illustration of aspects of a system according to an embodiment herein;

FIG. 7 is an illustration of aspects of a system according to an embodiment herein;

FIG. 8 is an illustration of aspects of a system according to an embodiment herein;

FIG. 9 is a flow diagram of a method according to an embodiment herein;

FIG. 10 is an illustration of articles of manufacture according to embodiments herein;

FIG. 11 is a schematic diagram of a hardware system according to embodiments herein.

FIG. 12 is a schematic diagram of a deployment system according to embodiments herein;

FIG. 13 is a schematic diagram of an integration system according to embodiments herein;

FIG. 14 is a schematic diagram of an on demand system according to embodiments herein;

FIG. 15 is a schematic diagram of a virtual private network system according to embodiments herein; and

FIG. 16 is a schematic diagram of a virtual private network system according to embodiments herein.

DETAILED DESCRIPTION

It will be readily understood that the embodiments herein, as generally described and illustrated in the figures herein, may be arranged and designed in a wide variety of different configurations in addition to the described embodiments. Thus, the following detailed description of the embodiments herein, as represented in the figures, is not intended to limit the scope of the invention, as defined by the appended claims, but is merely representative of selected embodiments. The following description is intended only by way of example, and simply illustrates certain selected embodiments herein.

The QA system described herein is not merely a document search system. One key difference between QA technology and simple document search technology is that document search typically takes a keyword query and merely returns a list of documents, ranked in order of relevance to the query. On the other hand, QA technology takes a question expressed in natural language, seeks to understand it in detail, and returns a precise answer to the question. A typical scenario is for an end user to enter a question in natural language form, much as if they were asking another person, and for the system to sift through vast amounts of potential evidence to return a ranked list of the most compelling, precise answers. These answers may also include summaries of their justifying or supporting evidence, allowing the user to quickly assess the evidence and select the correct answer.

Given sufficient search data, a document search can usually find documents or passages containing some keywords in common with the query, but lack the precision, depth, and understanding necessary to deliver correct answers with accurate confidence.

While an Internet web search engine has access to an immense source of information and can quickly find relevant web pages given a small number of query terms, such a search engine does not generate precise answers to the questions. A web search engine does not return answers; rather, it is designed to return a ranked list of web pages that the user may be trying to find.

A challenge for even finding the relevant documents is the issue of choosing the right set of keywords to retrieve those documents. Many questions contain information that is not critical for answering the question, but is provided for educational and/or entertainment purposes. Moreover, the question may use terms that are different from those used in evidence that may contain a correct answer. As a result, formulating an effective query that returns the relevant documents is a critical and non-trivial task.

In some cases, the question requires an answer that must be synthesized from multiple pieces of evidence from multiple source documents; that is, the answer is a list of items or a logical combination of two or more items. These answers do

not appear in any one place. Rather, the QA technology must synthesize the answer from independent sources to form a final answer.

Classic knowledge-based artificial intelligence approaches to QA systems try to logically prove an answer is correct from a logical encoding of the question and all the domain knowledge required to answer it. Such approaches are stymied by the prohibitive time and manual effort required to acquire massive volumes of knowledge and formally encode it as logical formulas accessible to computer algorithms, and the difficulty of understanding natural language questions well enough to exploit such formal encodings if available. Consequently, QA systems tend to falter in terms of breadth, but when they succeed, they are very precise.

Aspects of the various embodiments implement the use of failures in QA systems to motivate inquiries for common sense knowledge, and employ crowd sourcing to respond to those inquiries. These failures can be used in QA processes to determine what common-sense knowledge would be useful to have in order to improve answer retrieval. One aspect of the QA system is to be able to discover and pose follow-on inquiries to a user (or an external expert community) that, if answered, will improve the ability of the QA system to understand and evaluate supporting evidence for questions. Additionally, the acquired common-sense knowledge can be applied either off-line or during a live question answering session. Furthermore, specific QA process failures can be identified and used to specify specific gaps in the system's common-sense knowledge, which can then be filled. Moreover, specific user-interface elements for common-sense knowledge questions are provided and employed.

According to embodiments herein, when the QA system is unable to obtain an answer with a high degree of confidence, the QA system generates and poses follow-on inquiries in order to acquire information that can change the degree of confidence in an answer. The QA system solicits responses to the follow-on inquiries from external sources either in online mode, in which case the external sources' response to the follow-on inquiry will influence the QA system's answer to the current question, or in offline mode, where follow-on inquiries are dispatched to an external expert community at a later time.

It is often the case that the information need is not well captured by the QA system, as the question processing part may fail to classify the question properly or the information needed for extracting and generating the answer is not easily retrieved. In such cases, the QA system may reformulate the question or request assistance from an outside source.

The use of an external source of information to obtain information is known in the art as crowd sourcing. Crowd sourcing is a process that involves sending tasks to a distributed group of people. In crowd sourcing, the task or problem is typically sent to an undefined public rather than a specific body. Using failures of the QA system to motivate inquiries for common sense knowledge and employing crowd sourcing to respond to those inquiries, embodiments herein can evaluate agreement among the various respondents regarding responses to the inquiries in order to validate that the respondents are human and that the responses were not generated by an automated script.

In some embodiments, a question may be related to an interrogatory searching for specific information, such as, "What is the capital of New York?" In some embodiments, a question may be posed in the form of an answer or statement, such as done in the popular game show "Jeopardy!"

FIG. 1 is a flow diagram illustrating one non-limiting embodiment herein. According to FIG. 1, one embodiment

described herein enables a computer implemented method for enhancing the accuracy of a question-answer system in providing responses to questions inputted to the question-answer system.

The first step of the method, at **110**, involves providing the QA system. The QA system uses multiple question answering processes and evidence from a corpus of data to provide answers to questions. The details regarding processes and evidence for an example QA system that can be used in a preferred embodiment are described in co-pending U.S. patent application Ser. No. 12/152,411, incorporated herein by reference.

In the next step of the method, at **120**, a question is input to the QA system. The question may be in the form of a statement or an answer seeking an appropriate question.

Next, at **130**, a failure in a question answering process is determined. The QA system generates one or more candidate answers to the question with associated confidence scores based on results from scoring processes/algorithms for pieces of evidence extracted from a corpus of data. A failure may result either from an inability to generate a candidate answer with a confidence score above a threshold value or if the QA system cannot correctly interpret the question. Additionally, a failure may also result from an individual piece of evidence receiving a score below a threshold value.

At **140**, the failure is used to determine a missing piece of information. The missing piece of data/information may be data/information that would enable the QA system to improve a score for a piece of evidence, for example a passage, wherein the score for the piece of evidence is used in a confidence score for a candidate answer. The missing information may be any item of data, a fact, a syntactical relationship, a grammatical relationship, a logical rule, a taxonomy rule, a grammatical rule, or any other information that would increase a determined score for a piece of evidence that may support or refute a candidate answer to the question.

Next, at **150**, a follow-on inquiry is output to obtain the missing piece of information. The inquiry may be directed to outside sources that can include a variety of users in an expert community who may be human users or may be other electronic systems capable of providing a response, such as other QA systems. A follow-on inquiry may involve, for example, keyword matching, expansion of the original question, and/or a request for lexical semantic relationships. For example, the QA system might request a clarification of what sense a word is being used, or what type of information is being requested in the question.

At **160**, the QA system receives a response to the follow-on inquiry. The response is returned by a human user, expert community or other QA system. At **170**, the response to the follow-on inquiry is validated to confirm the missing piece of data. The validation may include validation that the response is supported by a threshold number of experts, humans, or QA systems.

Finally, at **180**, the obtained missing piece of data is added into the question-answer system. Again, the missing piece of data may comprise any item of data, a fact, a syntactical relationship, a grammatical relationship, a logical rule, a taxonomy rule, a grammatical rule, or any other information that would increase a determined score for a piece of evidence that may support or refute a candidate answer to the question. The missing piece of data may be input into the corpus, algorithm, process, logical rule, or any other location or combination thereof wherein the data may affect the resulting score for a piece of evidence.

In some embodiments, at least one step of the method is carried out using a computer device. In some embodiments,

the response to the follow-on inquiry may be evaluated to determine whether or not the response was provided by a human respondent.

Referring to FIG. 2, a general workflow of the QA system **210** is shown. The QA system **210** receives a question or questions **220** and searches a corpus of data **225** to extract pieces of evidence **230,240,250** that are evaluated by a number of algorithms in order to generate candidate answers **285** to the question **220**. The QA system **210** attempts to generate many candidate answers **285** to the question **220** presented. To do so, the QA system **210** considers a wide variety of information, e.g., texts of all kinds, web documents, reference books, novels, plays, encyclopedias, dictionaries, thesauri, textbooks, and technical reports, included in the corpus of data **225**.

The corpus of data **225** includes quantities of digitally encoded, structured and unstructured information, such as fragments of documents, passages, internal or external knowledge bases, natural language documents, corporate intranets, reference books, textbooks, technical reports, blogs, etc. These lists are intended to be non-limiting examples of sources of information for the QA system **210**. The vast majority of the corpus of data **225** comprises a wide variety of natural language text. Some of the data may be structured (formal knowledge bases) and some of the data may be semi-structured (tagged text) or unstructured (untagged text) in order to help interpret text and refine answers.

As used herein, 'pieces of evidence' refers to the specific information or data that the QA system **210** evaluates to refute or support candidate answers **285** to the question **220**. The QA system **210** draws pieces of evidence from the corpus of data **225**. The term piece of evidence is not intended to limit the length of evidence used and may include a document, passage from a document, data, meta-data, entry in a database or any other information capable of contributing to a confidence score for a candidate answer. Additionally, as used herein, 'logical rules' refer to the lexical, grammatical, and/or semantic relationships that the QA system **210** uses to generate and score candidate answers **285** to the question **220** based on the pieces of evidence. The logical rules enable the QA system **210** to draw an inference from the pieces of evidence in order to evaluate the candidate answers **285**.

In trying to generate candidate answers **285**, the QA system **210** extracts pieces of evidence from the corpus of data **225**. The pieces of evidence, based on the question, can be fit into three categories: good evidence **230**, bad evidence **240**, and marginal evidence **250**. The QA system **210** compares the question **220** to the data in the corpus of data **225** and assigns an evidence score to the extracted evidence based on how well the evidence matches the question **220**. Sometimes there is data in the corpus of data **225** that clearly and directly relates to the input question. This is referred to as "good" evidence. Good evidence **230** is a piece of evidence with a high evidence score; that is, an evidence score above a previously established evidence threshold value such that it contributes significantly to a confidence score for a candidate answer. Conversely, bad evidence **240** is a piece of evidence with such low relevance that it does not contribute to a confidence score for a candidate answer. A QA system **210** may not extract bad evidence from the corpus, or may extract bad evidence in order to evaluate and determine outright that this evidence will not provide a candidate answer with any confidence. Lastly, there are also pieces of evidence that the QA system **210** is able to identify as relevant to the question and may produce an evidence score that contributes significantly to a confidence score for a candidate answer if a missing link or missing piece of information is provided. A missing link or

missing piece of information could be described as all the marginal evidence, or a subset of the marginal evidence (either all relatively high scoring marginal evidence, all relatively low scoring marginal evidence, or some middle range of scores for the marginal evidence (based on predetermined score parameters)). Thus, if a piece of marginal evidence falls within some predetermined score range, it would be considered a piece of missing information, for which additional information will be obtained. In other words, the marginal evidence **250** is evidence that the QA system **210** can use in generating candidate answers **285**, but is lacking in some manner, such that the score for the piece of evidence does not contribute significantly to the overall confidence score for a candidate answer **285**. A low confidence score for all candidate answers **285** indicates that the QA system **210** should obtain additional information or assistance to further develop the confidence scores of the candidate answers **285**. In other words, low confidence scores may serve as a trigger for the QA system **210** to evaluate whether pieces of evidence **230**, **240**, **250** for a candidate answer **285** are marginal evidence **250** such that a missing information or missing link may be identified.

The QA system **210** uses the good evidence **230** and the marginal evidence **250** to generate candidate answers **285**. Each candidate answer **285** is assigned a confidence score, which is an aggregation of the evidence scores. A candidate answer **285** with a low confidence scores; that is, confidence scored below a previously established confidence threshold value indicates that the QA system **210** should obtain additional information or assistance to further develop the confidence scores of the candidate answers **285**. In other words, bad evidence does not contribute to the candidate answers **285**, marginal evidence contributes only partially to the candidate answers **285**, and good evidence provides high confidence in the candidate answers **285**.

For a candidate answer generated, the QA system **210** evaluates a self-assessed probability that the candidate answer **285** is a correct answer to the question **220**, or in other words, the QA system **210** generates a confidence score for a candidate answer **285**. In order to generate candidate answers **285**, the QA system **210** analyzes the question to figure out what is being asked. It also analyzes the corpus of data **225** based on the question **220** in order to extract candidate answers **285**, and computes a reliable confidence in light of whatever supporting or refuting pieces of evidence **230**, **240**, **250** it finds. To obtain confidence in a candidate answer **285**, the QA system **210** attempts to justify the answer by multiple sources in the corpus of data **225**. The degree to which the probability that the candidate answer **285** is a correct answer to the question **220** correlates to the confidence score for the candidate answer **285**. The QA system **210** can then rank the candidate answers **285** based on the confidence scores.

The QA system **210** generates follow-on inquiries **260** to supplement the marginal evidence **250** by identifying an aspect of the piece of evidence that is lacking, in other words, identifying the missing piece of information or missing link. The QA system **210** uses follow-on inquiries **260** to clarify the aspect of the evidence that the QA system **210** failed to understand. According to an embodiment herein, the follow-on inquiries **260** are sent to an external source **272**, which may include a community of potential respondents **270**, which may be on an electronic network, such as the Internet **275**. After the responses **280** to the follow-on inquiries have been evaluated, the validated responses may then be fed back into the QA system **210** as data, a fact, a syntactical relationship, a grammatical relationship, a logical rule, a taxonomy rule, a grammatical rule, or any other information that would

increase a determined score for a piece of evidence that may support or refute a candidate answer to the question. The potential respondents **270** may be from an external expert community who may be human users or other electronic systems capable of providing a response, such as other QA systems. The responses **280** to the follow-on inquiries may include the missing information or missing link necessary to make the marginal evidence **250** useful in order to improve the confidence scores of one or more candidate answers **285**. Such missing information can be used to produce improved candidate answers **285** to the inputted question or can be used to produce subsequent follow-on inquiries **260**. In either case, the information can be used to enhance the internal knowledge base of the QA system **210** and improve the confidence in the candidate answers **285**. The missing information included in the responses **280** to the follow-on inquiries can also help repair flaws in the question answering processes used by the QA system **210** by incorporating the information of the response, and the manner that provided the response, into the question answering processes.

In some embodiments herein, the QA system **210** evaluates the responses **280** to the follow-on inquiries to determine agreement among the various respondents regarding the responses **280** to the follow-on inquiries, in order to validate that the respondents are human and that the responses were not generated by an automated script.

Inputs to the QA system **210** may include questions **220**, the evidence extracted from the corpus of data **225**, and responses **280** to follow-on inquiries. Outputs from the QA system **210** include candidate answers **285** and follow-on inquiries **260**. This provides low cost, continuous knowledge acquisition.

FIG. 3 shows one approach to identifying marginal evidence **250**. The QA system **210** receives a question **320** and parses the question into a collection of elements **340**. The QA system **210** searches the corpus of data **225** for a piece of evidence, for example a passage, such as **330**, to support or refute a candidate answer **285** to the question. The QA system **210** parses the passage **330** into a collection of elements **350**. The QA system **210** analyzes the passage **330** by seeking to match, or associate in some manner, the elements **340** in the question **320** and the elements **350** in the passage **330** based on lexical, grammatical, and/or semantic relationships. The QA system **210** may, for example, attempt to match the various individual elements, such as **345**, **355** with corresponding semantic roles in the question **320** and the passage **330**. The QA system **210** then classifies the passage as good, bad, or marginal as indicated at **360**. As used herein, 'marginal' means that some part of the passage **330** may be relevant, but it needs additional information to improve a score for the piece of evidence based on a scoring algorithm or process. This additional information can be labeled as a "missing link" or "missing information" that, when connected or combined with the other information in the passage **330**, enables the QA system **210** to further develop the confidence score of the candidate answers **285**.

In the approach illustrated in FIG. 3, the inputs to the QA system **210** include the question **320** and the passage **330** that is used for comparison to the question **320**. Outputs from the QA system **210** include classification of the passage **330** into good, bad, and marginal matches to the question **320**, or simply identification of the marginal evidence without requiring identification of the good or bad evidence. The marginal evidence includes missing information thus follow-on inquiries **260** can be generated in order provide the missing information from an external source.

The QA system 210 may employ a variety of specific mechanisms for finding marginal evidence. Aligning questions to sample passages is one of the ways that the QA system 210 uses evidence to evaluate candidate answers 285. Given an answer key, it can be used to drive the search for near misses. In general, there are many kinds of “near miss” failures that the QA system 210 could encounter, for example:

Answer typing: the system could have determined that the answer had the correct type if only it had known that two types were equivalent

Keyword search: the system could have found a relevant passage using a keyword search if only it had expanded some keyword in the clue; e.g., by synonyms, antonyms, related terms, etc.

Relation detection for knowledge-base lookup: the system could have looked up the answer in a knowledge base if only it had figured out what relationship the question was asking for.

Each kind of near miss can imply some kind of semantic inquiry.

In some embodiments, mismatched elements, such as 345, 355, may actually be equivalent or have a counterpart as a logically necessary consequence. A follow-on inquiry 260 can be used to determine whether the mismatched elements are equivalent. Such follow-on inquiry 260 can use the equivalence to decide if the passage 330 aligns with the question 320, and can also be used to resolve similar alignment problems in future questions.

The QA system 210 analyzes lexical, grammatical, and semantic relationships of the words within each question 320; therefore, the QA system 210 breaks each question 320 down into many separate elements 340 (including keywords, equivalent keywords, logical relationships, restrictions, subject categories, etc.). For example, the good evidence 230 matches a sufficient number of the elements to provide a high score for the piece of evidence; that is, a score above a threshold value that contributes significantly to a confidence score for a candidate answer 285. For example, the marginal evidence 250 matches some of the elements of the question 320, but not all such elements, resulting in a score for the piece of evidence below a previously established threshold value. The systems and methods herein identify the elements of the question 320 that the marginal evidence elements were not able to match. Thus, the systems and methods herein identify the portions of the question 320 that were not provided by the marginal evidence 250 and generate follow-on inquiries 260 to request the external sources 272 to supply information to clarify the portions of the question 320 that were not provided by the marginal evidence 250. In some embodiments, the follow-on inquiries 260 may use a natural language format.

In other words, the systems and methods herein automatically identify the missing link that, if known, would increase the confidence score for one or more candidate answers 285. This missing link can be an item of data, a fact, a syntactical relationship, a grammatical relationship, a logical rule, a taxonomy rule, a grammatical rule, etc. By generating the follow-on inquiry, systems and methods herein can receive a number of responses (from, for example, a crowd sourcing environment). These responses can be filtered to make sure that the responses are from actual humans (as opposed to computerized answering systems), for example, based on responses to supply missing information that a machine would have trouble answering. Additionally, the responses can be filtered to make sure that the responses are above a certain level of accuracy (based upon the reputation of the responder, correlation with the known correct answer, agreement of the responses with one another, etc.). Then, once the

responses are filtered, the high-quality, human-based responses can be utilized to increase the knowledge base of the question answer system, and/or to generate additional rules to help the question answer system score and rank the candidate answers 285.

FIG. 4 demonstrates a practical application of the approach to identify marginal evidence illustrated in FIG. 3. In FIG. 4, the QA system 210 compares the question 420 to the passage 430. In the example of FIG. 4, the question 420 is “Eleazar killed one of these huge beasts.” In other words, the question 420 is looking for an answer as to what type of beast was killed by Eleazar. In other aspects of the QA system 210 not shown, the piece of evidence, or passage 430, has been identified as relevant to answering the question 420. In performing the analysis between the passage 430 and the question 420, the QA system 210 evaluates the text for lexical, grammatical, and/or semantic relationships among the elements 440 of the question 420 and the elements 450 of the passage 430. In the example of FIG. 4, the QA system 210 can match some elements, such as Eleazar 445, to conclude that the passage is of some relevance. However, the QA system 210 may not be able to determine that “stabbed” 457 and “killed” 447 have a logical relationship. This failure of the QA system 210 reveals the missing information or missing link 460. In this example, the missing link is to resolve whether “killed” 447 is logically related to “stabbed” 457. Therefore, in this example, the systems and methods herein could generate a follow-on inquiry 260 that asks for a crowd source response to the question “does stabbing imply killing?” If the response 280 to this follow-on inquiry is to the affirmative, the systems and methods herein can automatically learn that the two words may be synonyms. According to an embodiment herein, the system concludes that the passage 430 would receive a higher score and contribute more significantly to the confidence score for a candidate answer 285 of “elephant” if “killed” is logically related to “stabbed.”

According to an embodiment herein, FIG. 5 provides examples of some specific examples of follow-on inquiries 260. As illustrated in Example 1 of FIG. 5, natural language follow-on inquiries 510 are actual natural language questions generated by the QA system 210 that can be supplied directly to external sources 272, such as another QA system, another information source, or human respondents 270 on an electronic network (such as the Internet 275), so that the QA system 210 can acquire the missing information needed to better respond to the question 220.

As illustrated in Example 2 of FIG. 5, given the context of the question and passage, the follow-on inquiries can be multiple-choice taxonomies using text and/or diagrams 515. As would be known in the relevant art, synonyms are different words with similar or almost identical meanings, a hyponym is a word or phrase whose semantic field is included within that of another word; for example: rose is a hyponym of flower, and a hypernym is a word with a general meaning that has basically the same meaning of a more specific word; for example, the word “color” is a hypernym for red, green, blue, brown, etc.

By acquiring responses in a crowd sourcing manner, the systems and methods herein allow the QA system 210 to answer questions more similarly to a human expert and to be customized for different localities. More specifically (with reference to answering the question of FIG. 4), purely logical rules may require that the correct answer in the multiple-choice example 515 be selection “D” because “stab” and “kill” sometimes overlap, but not always (surely logically eliminating answers A-C). However, depending upon local speech customs, “killing” and “stabbing” can be considered

synonyms. By automatically learning that “killing” and “stabbing” are considered synonyms in a sufficient number of situations, the systems and methods herein allow the QA system **210** to answer questions with higher confidence, and provides more flexibility in its word definitions and grammatical/logical rules. The ability of the systems and methods herein to automatically learn of human peculiarities by utilizing crowd sourced responses allows the QA system **210** to respond more consistently as a human would when generating answers to questions. This further underscores the need to filter out non-human responses **280** to the follow-on inquiry because non-human responses may give the correct, albeit purely logical response, if supplied by a machine-based answering system. The QA system **210** herein is able to receive information a human would consider to be a correct response (even if such a response might not be the most correct answer in a purely logical sense).

According to embodiments herein, the QA system **210** can receive a number of responses, for example, from a crowd sourcing environment or an external expert community. These responses can be filtered to make sure that the responses are from actual humans, as opposed to computerized answering systems, and that the responses are above a certain level of accuracy either based upon the reputation of the responder, correlation with the known correct answer, or agreement of the responses with one another.

An additional example of using follow-on inquiries is described below:

Question: This star of “Emmerdale Farm” also plays Elvis in “If You Ever Get to Heaven.”

In this example, the question is looking for a “star”; however, the QA system **210** may not know if that means a star in the sky or a star on the stage. An appropriate answer to the question is “Samuel Kane”. It is possible that a first source from the corpus of data identifies Samuel Kane as an actor. However, the QA system may not be able to determine that Samuel Kane is a “star.” As such, the QA system has identified missing information or a missing link regarding the lexical relationship between an actor and a star. In an embodiment herein, the QA system **210** may direct some follow-on inquiries. Some examples of follow-on inquiries may include:

Is every actor a star?

Is every star an actor?

Are most actors also stars?

Are many actors also stars?

If I know that some entity is a star, would I suspect that this entity is an actor?

If I know that some entity is an actor, would I suspect that this entity is a star?

If I know that some entity is an actor, can I be certain that this entity is a star?

Using follow-on inquiries, the QA system **210** tries to determine more about being an actor and being a star. By evaluating the responses to the follow-on inquiries, the QA system **210** determines that this question is asking for a “star.” If the QA system **210** finds a possible answer that is an actor, can it be certain that this answer is a star? Conversely, if the QA system **210** finds a possible answer that is an actor, should it suspect that this answer is a star? According to one resource, the QA system **210** has found a possible answer: Samuel Kane. Based on that information, the QA system **210** believes that Samuel Kane is an actor. Should the QA system **210** conclude that Samuel Kane is a star because it believes that Samuel Kane is an actor? Should the QA system **210** suspect that Samuel Kane is a star because it believes that Samuel Kane is an actor?

An additional example herein employs question expansion to generate follow-on inquiries, as follows:

Question: “This television quiz show host from Mexico was born in the last few days of the 60’s.”

In this example, the question is looking for several specific elements. An appropriate answer to the question is “Marco Antonio Regil.” According to an embodiment herein, the QA system **210** uses the question text to build a keyword search question. The QA system **210** considers many possible expansions of the question; e.g., show is related to program, born is related to birth, television is related to TV, etc. There is a relevant piece of evidence or passage that states, “Marco Antonio Regil (born Dec. 27, 1969) is a Mexican television personality and game show host, born in Tijuana, Mexico.” An embodiment herein can find this relevant passage since it contains the correct answer and some keywords from the clue. The QA system **210** can then look for expansions to the original question that would have produced this passage. The QA system **210** finds several expansions: quiz show is related to game show, Mexico is related to Mexican, 60’s is related to 1969. These could be used as training data to train a more effective question expander if they are actually implied by the clue. Thus, this embodiment can enable the following follow-on inquiries:

If a question is requesting information about a quiz show, can I be certain that information about a game show is relevant?

If a question is requesting information about a quiz show, should I be suspect that information about a game show is relevant?

If a question is requesting information about the 60’s, can I be certain that information about 1969 is relevant?

If a question is requesting information about the 60’s, should I be certain that information about a 1969 is relevant?

Additionally, the QA system **210** can present follow-on inquiries to potential respondents **270** (internet users, other QA system, etc.) requesting potential respondents **270** to rank the terms of the question from most relevant to the topic of this question to least relevant:

a) game show

b) birth

c) Mexican

d) program

e) TV

f) 1969

The QA system **210** can present follow-on inquiries to potential respondents **270** (internet users, other QA system, etc.) requesting potential respondents **270** to rate the terms on a number scale based on relevance to the topic of this question, for example rate from 1-10 in term of relevance the following terms:

a) game show

b) birth

c) Mexican

d) program

e) TV

f) 1969

Another embodiment herein employs relation detection for knowledge-based lookup to generate follow-on inquiries. An example of relation detection is described below:

Question: “She appeared on the big screen as Rachel Powell in the Hollywood blockbuster K-PAX.”

In this example, the question is looking for a specific person’s name. An appropriate answer to the question is “Mary McCormack.” According to an embodiment herein, the QA system **210** may use the question text and a structured knowl-

edge base to arrive at the correct answer. The structured knowledge base may have an entry labeled Role_In_Movie (“Mary McCormack”, “Rachel Powell,” “K-PAX”). Thus, the QA system 210 would have gotten the answer correct had it concluded that the question was asking for the specific variable that satisfies the following structured question: played_Role_In_Movie (???, “Rachel Powell,” “K-PAX”). However, the system does not have a detection pattern for Played_Role_In_Movie that matches this clue. It does have a pattern for Played_Role_In_Movie (X,Y,Z) of the form “X played Y in Z”. Thus, it knows that if the text had been “She played Rachel Powell in K-PAX,” then it would have detected this relation and gotten the question right; and as such, it could present a context like this to the respondents 270. However, the QA system 210 did not fully understand the following question: “She appeared on the big screen as Rachel Powell in the Hollywood blockbuster K-PAX.” It thinks it may be able to simplify this question into the following question: “She played Rachel Powell in K-PAX.” It could then present one or more of the following questions given this context: Is this a valid simplification? Can you imagine a correct answer to the simplified form that is not a correct answer to the original question? Can you imagine a correct answer to the original question that is not a correct answer to the simplified form? The original question uses the word “Hollywood” but the simplified form does not. Is “Hollywood” implied in the simplified form? The original question uses the word “Hollywood” but the simplified form does not. Is “Hollywood” a very important concept for answering this question? Please rank the following terms from the original clue that are not in the simplified form as most important to least important in terms of what was lost in simplifying:

- a) appeared
- b) big
- c) screen
- d) Hollywood
- e) blockbuster

To compute the confidence scores for the candidate answers, the QA system 210 uses several criteria: e.g., is the candidate answer satisfying similar lexical and semantic relations (e.g. for a query about an actress starring in a movie, is the candidate answer a female, and does the candidate answer satisfy actor-in-movie relation?); how well does the candidate answer and the query align; how well do the terms match, and do the terms exist in similar order.

The language and knowledge processing infrastructure of the QA system 210 combines statistical and heuristic techniques to assess its own knowledge and produce its best candidate answer 285 with an accurate confidence—a measure of the likelihood it is correct based on a self-assessment of its sources, inference methods, and prior performance. Unfortunately, some of the respondents 270 on the Internet 275 may provide incorrect or misleading result. As noted above, the systems and methods herein can utilize a crowd sourcing technique to obtain responses to the follow-on inquiries 260. Crowd sourcing involves asking external sources 272, such as an expert community, which can be many human Internet/computer users, to provide individual responses to the follow-on inquiries. Understandably, some of the responses from some respondents 270 will be incorrect. However, the “correctness” of the response is a relative statement because, as shown above in the examples in FIG. 5, some semantic and grammatical human ways of asking and answering questions may not be purely logically “correct.” Therefore, what is considered a strong synonym by humans may not be considered a synonym under purely logical rules. Thus, the systems and methods herein learn what the majority

of the respondents 270 consider to be the “correct” responses to the follow-on inquiries and use such as a standard of what is “correct” in order to become more humanlike.

Further, what is considered to be the “correct” response to a follow-on inquiry 260 can vary from geographic place to place, depending upon local language usage of humans in the different areas. Therefore, the systems and methods herein not only learn from the majority of all respondents as to what is “correct”, but also learn what the “correct” response is for each of the different geographic areas in which the QA system 210 may be utilized. That is, what may be considered a very close synonym is North America may not be considered a synonym at all in Ireland despite sharing the common English language. The systems and methods herein, therefore, learn what humans consider to be the correct answer component on many different levels.

According to embodiments herein, the QA system 210 can employ many alternative approaches to validating responses, three examples of which are illustrated in FIGS. 6-8. More specifically, FIGS. 6, 7, and 8 illustrate verification schemes that may be employed to determine the accuracy of the information submitted to the QA system 210.

According to FIG. 6, one embodiment herein permits only outside sources with demonstrated accuracy in providing previous information 645, such as correct answers 285 to previous questions, to provide answers to follow-on inquiries 260. The QA system 210 can fill-in and classify marginal evidence, and supply answers generally. In the embodiment of FIG. 6, the same follow-on inquiry 260 is given to multiple potential respondents 270. If several respondents 270 provide the same result, the QA system 210 can generally trust the answer.

According to FIG. 7, one embodiment herein permits outside sources to supply answers that match with the results the QA system 210 has determined regarding all the good and bad evidence 775. Using results that the QA system 210 has already marked as “good” 750 or “bad” 760, the QA system 210 can generally trust the respondents 270 that agree with the QA system 210 on the good evidence and bad evidence 775. Since the accuracy of the QA system 210 regarding the classifications of the good evidence 750 and the bad evidence 760 is beyond dispute, it can be implied that respondents 270 that agree with these assessments are more likely than not to be accurate in order to provide responses to follow-on inquiries 780 that match with its own assessments, can fill-in and classify marginal evidence, and supply responses generally.

According to FIG. 8, one embodiment herein permits specifically identified outside sources to supply answers 885 if they demonstrate a sufficient amount of accuracy in relation to a known trusted source 880, such as a human expert. According to this embodiment, a trusted source 880 provides a response for several examples. The QA system 210 can generally trust respondents 270 that often agree with the trusted source 880 on those examples.

Accuracy in any of the embodiments can be determined by tracking the success rate of the outside source or sources 272 in relation to the metrics and comparisons indicated in the above embodiments. This can be done by keeping track of responses in relation to the above embodiments for definite duration or an infinite duration, and it can be done internally by the QA system 210, or by another tracking system or method known in the art.

Furthermore, the embodiments disclosed in FIGS. 6, 7, and 8 can be employed individually, or they can be combined in any manner. In some embodiments, the verification techniques illustrated in FIGS. 6, 7, and 8 can be used in hybrid verification schemes by mixing the features of portions of

each method. For example, the verification scheme shown in FIG. 8 can compare the determinations of the reliable trusted source 880 on good and bad evidence 775 to those of the other respondents 270 to determine which outside source will be permitted to respond to the QA system 210. Similarly, the verification scheme can be a combination metric of answers provided to questions, answers to follow-on inquiries, and/or classifications regarding evidence provided by the reliable outside source in relation to the answers provided to questions, answers to follow-on inquiries, and/or classifications regarding evidence provided by the other outside sources. Furthermore, the verification scheme can combine any of the methods in any manner in order to generate a hybrid metric that permits the outside source to respond to the QA system 210.

All of the above schemes and methods enable a QA system to expand its knowledge base as can be contained in a database, correct its processes for future questions using external sources, such as human users on a network or other QA systems, and generally increase its accuracy. As should be apparent, the external or outside sources referenced throughout this application, whether explicitly stated or not, include, but are not limited to, sources accessible to the QA system 210 such as human sources, electronic sources, such as another QA system, or any other source accessible by the QA system 210 that is capable to facilitate responses to the QA system 210 according to the principles embodied in the present disclosure.

The QA system 210 described herein is an application of advanced natural language processing, information retrieval, knowledge representation and reasoning, and machine learning technologies to the field of open-domain question answering using technology for hypothesis generation, massive evidence gathering, analysis, and scoring. In addition to identifying an answer 285, the QA system 210 described herein determines a confidence in its answer 285. If that confidence is low, the system realizes that it may not understand the question. For example, there may be a pun or some syntax that the QA system 210 does not understand. On the other hand, if the confidence is high, the QA system 210 knows it likely understood the question and stands a good chance of getting that question right. Having sufficient information, the QA system 210 can always find documents or passages containing some keywords in common with the question but may lack the precision, depth, and understanding necessary to deliver correct answers with accurate confidence.

An additional embodiment herein uses the failures of the QA system to generate questions to determine whether the outside responding source is a human source. As described above, the failures of the QA system motivate follow-on inquiries to employ common sense knowledge and crowd sourcing techniques in order to respond to those follow-on inquiries. According to another embodiment, the QA system can evaluate the nature of the responses received and determine that a respondent is human.

FIG. 9 is a flow diagram illustrating one non-limiting embodiment herein. According to FIG. 9, one embodiment described herein enables a computer implemented method for validating that a user is human.

The first step of the method, at 910, involves providing a question answering system. The QA system uses question answering processes and evidence passages to provide responses to questions.

In the next step of the method, at 920, a question is input to the QA system. The question may be in the form of a statement or an answer seeking an appropriate question.

Next, at 930, a failure in a question answering process is determined. The QA system generates one or more candidate answers to the question with associated confidence scores based on results from scoring processes/algorithms for pieces of evidence extracted from a corpus of data. A failure may result either from an inability to generate a candidate answer with a confidence score above a threshold value or if the QA system cannot correctly interpret the question. Additionally, a failure may also result from an individual piece of evidence receiving a score below a threshold value.

At 940, the failure is used to determine a missing piece of information. The missing piece of data/information may be data/information that would enable the QA system to improve a score for a piece of evidence, for example a passage, wherein the score for the piece of evidence is used in a confidence score for a candidate answer. The missing information may be any item of data, a fact, a syntactical relationship, a grammatical relationship, a logical rule, a taxonomy rule, a grammatical rule, or any other information that would increase a determined score for a piece of evidence that may support or refute a candidate answer to the question.

Next, at 950, an inquiry is output to outside sources in order to obtain the missing piece of data. The inquiry may be directed to outside sources that can include a variety of users who can be human users or electronic systems capable of providing a response, such as other QA systems. A follow-on inquiry may involve, for example, keyword matching, expansion of the original question, and/or a request for lexical semantic relationships. For example, the QA system might request a clarification of what sense a word is being used, or what type of information is being requested. In some embodiments, the follow-on inquiry may ask a question for which the answer is known by the QA system. In some embodiments, the follow-on inquiry may ask a question for which the answer is not known by the QA system.

At 960, after receiving a response from the responding source, the QA system evaluates the nature of the response to determine whether the outside or external responding source is human 970.

The QA system 210 can do this by keeping track of how other humans have responded to questions or follow-on inquiries, and it can compare a response from an external source to the tracked responses. Responses that agree with responses from other humans can indicate that a respondent is human. These can be the same, similar, or different follow-on inquiries and questions. If the response matches well, the QA system can determine that the respondent was human. Afterwards, the QA system may determine the answers to questions for which the answer is not already known based on consensus among users who were previously known to be human.

Another method that the QA system can use to determine if the responding source is human is to phrase the question or follow-on inquiry in natural language fashion. The QA system can then assess the accuracy of the response and/or determine whether or not the response itself is in natural language format. To determine if a response is in natural language format the system can perform several types of analyses. In one non-limiting example, the QA system 210 can perform a syntactic analysis by parsing the answer to determine if the answer conforms to syntactic rules of a language, for example, English. Similarly, the QA system 210 can check to see if the answer is semantically similar by comparing the predicate argument structure to a collection of human produced text. For example, the phrase "A quasar can star in a horror movie" is syntactically well formed but the predicate "star" never occurs together with "quasar" and "honor

movie.” This can be checked in a lexical database such as Prismatic, as described in J. Fan et al., “Prismatic: Introducing Knowledge from a Large Scale Lexical Relation Resource,” Proceedings of the NAACL HLT 2010 First International Workshop on Formalisms and Methodology by Reading, pages 122-127, June 2010. In some embodiments, the QA system can use general questions, natural language or otherwise, that humans have a known particular way of responding. Using such general questions, the QA system can determine whether the responding source is human. These descriptions are merely exemplary in nature. It is contemplated that those skilled in the art can conceive and apply additional methods and techniques that will achieve the same result in accordance with the present disclosure and its principles.

The present invention applies to all QA systems now known and as may be developed in the art, and can be implemented by computer or electronic devices now known (or that may become known) in the art.

According to another embodiment herein, a question answering (QA) system includes a processor, an evidence analysis module connected to the processor, a first interface connected to the processor and a second interface connected to the processor, and a corpus of data connected to the evidence analysis module. The QA system is the type capable of producing evidence in support of a candidate answer. A first question to be answered by the QA system is provided through the first interface. The processor creates a collection of candidate answers to the first question from the corpus of data. Each candidate answer has supporting evidence and a score generated by the processor. The evidence analysis module produces a second question (i.e. follow-up inquiry) based on the supporting evidence. The processor presents the second question through the second interface to one or more users. The processor receives at least one response or knowledge item from one or more users through the second interface. The processor validates the at least one response and determines whether a user is human based on comparing the at least one response or knowledge item with other answers provided to the second question. In some embodiments, the processor inputs validated responses into the corpus of data to improve the ability of the QA system to answer questions. The new response or knowledge item changes the supporting evidence for the first question and the scores for the candidate answers. Based on the changed evidence, the QA system generates an answer to the first question through the first interface. The system can be scaled up or down depending on different application requirements.

According to a further embodiment herein, an article of manufacture is provided that includes a computer readable medium having computer readable instructions embodied therein for performing the steps of the computer implemented methods, including but not limited to the methods illustrated in FIGS. 1 and 9. Any combination of one or more computer readable non-transitory medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. The non-transitory computer storage medium stores instructions, and a processor executes the instructions to perform the methods described herein. 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. FIG. 10 illustrates exemplary articles of manufacture, such as, a magnetic storage device 560, a portable compact disc read-only memory (CD-ROM) 570, and a “plug-and-play” memory device 580, like a USB flash drive. Any of

these devices may have computer readable instructions for carrying out the steps of the methods described above with reference to FIGS. 1 and 9.

Aspects of the present disclosure are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems) and computer program products according to various embodiments. It will be understood that each block of the flowchart illustrations and/or two-dimensional block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer program instructions. The 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 execute 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.

The computer program instructions may 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.

Furthermore, 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 execute 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.

In case of implementing the embodiments herein by software and/or firmware, a program constituting the software may be installed into a computer with dedicated hardware, from a storage medium or a network, and the computer is capable of performing various functions if with various programs installed therein.

A representative hardware environment for practicing the embodiments herein is depicted in FIG. 11. This schematic drawing illustrates a hardware configuration of an information handling/computer system in accordance with the embodiments herein. The system comprises at least one processor or central processing unit (CPU) 610. The CPUs 610 are interconnected via system bus 612 to various devices such as a random access memory (RAM) 614, read-only memory (ROM) 616, and an input/output (I/O) adapter 618. The I/O adapter 618 can connect to peripheral devices, such as disk units 611 and tape drives 613, or other program storage devices that are readable by the system. The system can read the inventive instructions on the program storage devices and follow these instructions to execute the methodology of the embodiments herein.

In FIG. 11, CPUs 610 perform various processing based on a program stored in a Read Only Memory (ROM) 616 or a program loaded from a peripheral device, such as disk units 611 and tape drives 613 to a Random Access Memory (RAM) 614. In the RAM 614, required data when the CPU 610 performs the various processing or the like is also stored as necessary. The CPU 610, the ROM 616, and the RAM 614 are connected to one another via a bus 612. An input/output adapter 618 is also connected to the bus 612 to provide an input/output interface, as necessary. A removable medium,

such as a magnetic disk, an optical disk, a magneto-optical disk, a semiconductor memory, or the like, is installed on the peripheral device, as necessary, so that a computer program read therefrom may be installed into the RAM 614, as necessary.

The system further includes a user interface adapter 619 that connects a keyboard 615, mouse 617, speaker 624, microphone 622, and/or other user interface devices such as a touch screen device (not shown) to the bus 612 to gather user input. Additionally, a communication adapter 620 including a network interface card such as a LAN card, a modem, or the like connects the bus 612 to a data processing network 625. The communication adapter 620 performs communication processing via a network such as the Internet. A display adapter 621 connects the bus 612 to a display device 623, which may be embodied as an output device such as a monitor (such as a Cathode Ray Tube (CRT), a Liquid Crystal Display (LCD), or the like), printer, or transmitter, for example.

In the case where the above-described series of processing is implemented with software, the program that constitutes the software may be installed from a network such as the Internet or a storage medium such as the removable medium.

Those skilled in the art would appreciate that the storage medium is not limited to the peripheral device having the program stored therein as illustrated in FIG. 11, which is distributed separately from the device for providing the program to the user. Examples of a removable medium include a magnetic disk (including a floppy disk), an optical disk (including a Compact Disk-Read Only Memory (CD-ROM) and a Digital Versatile Disk (DVD)), a magneto-optical disk (including a Mini-Disk (MD) (registered trademark)), and a semiconductor memory. Alternatively, the storage medium may be the ROM 616, a hard disk contained in the storage section of the disk units 611, or the like, which has the program stored therein and is distributed to the user together with the device that contains them.

As will be appreciated by one skilled in the art, aspects of the systems and methods herein may be embodied as a system, method, or computer program product. Accordingly, aspects of the present disclosure 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 disclosure 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.

Any combination of one or more computer readable non-transitory medium(s) may be utilized. The computer readable medium may be a computer readable signal medium or a computer readable storage medium. The non-transitory computer storage medium stores instructions, and a processor executes the instructions to perform the methods described herein. 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 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 magnetic storage device 560 (FIG. 10), a portable compact disc read-only memory (CD-ROM) 570, an

optical storage device, a "plug-and-play" memory device 580, like a USB flash drive, 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 an instruction execution system, apparatus, or device.

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 an instruction execution system, apparatus, or device.

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.

Computer program code for carrying out operations for aspects of the present disclosure 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 execute 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).

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 herein. 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 executable instructions for implementing the specified logical function(s). 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 executed substantially concurrently, or the blocks may sometimes be executed 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.

Deployment types include loading directly in the client, server, and proxy computers via loading a storage medium such as a CD, DVD, etc. The process software may also be automatically or semi-automatically deployed into a computer system by sending the process software to a central server or a group of central servers. The process software is then downloaded into the client computers that will execute the process software. The process software is sent directly to the client system via e-mail. The process software is then either detached to a directory or loaded into a directory by a

button on the e-mail that executes a program that detaches the process software into a directory. Alternatively, the process software is sent directly to a directory on the client computer hard drive. When there are proxy servers, the process will select the proxy server code, determine on which computers to place the proxy servers' code, transmit the proxy server code, and then install the proxy server code on the proxy computer. The process software will be transmitted to the proxy server, and then stored on the proxy server.

While it is understood that the process software may be deployed by manually loading directly in the client, server, and proxy computers via loading a storage medium such as a CD, DVD, etc., the process software may also be automatically or semi-automatically deployed into a computer system by sending the process software to a central server or a group of central servers. The process software is then downloaded into the client computers that will execute the process software. Alternatively, the process software is sent directly to the client system via e-mail. The process software is then either detached to a directory or loaded into a directory by a button on the e-mail that executes a program that detaches the process software into a directory. Another alternative is to send the process software directly to a directory on the client computer hard drive. When there are proxy servers, the process will select the proxy server code, determine on which computers to place the proxy servers' code, transmit the proxy server code, and then install the proxy server code on the proxy computer. The process software will be transmitted to the proxy server, and then stored on the proxy server.

In FIG. 12, step 700 begins the deployment of the process software. The first thing is to determine if there are any programs that will reside on a server or servers when the process software is executed 701. If this is the case, then the servers that will contain the executables are identified 809. The process software for the server or servers is transferred directly to the servers' storage via FTP or some other protocol, or by copying through the use of a shared file system 810. The process software is then installed on the servers 811.

Next, a determination is made on whether the process software is to be deployed by having users access the process software on a server or servers 702. If the users are to access the process software on servers, then the server addresses that will store the process software are identified 703.

A determination is made if a proxy server is to be built 800 to store the process software. A proxy server is a server that sits between a client application, such as a Web browser, and a real server. It intercepts all requests to the real server to see if it can fulfill the requests itself. If not, it forwards the request to the real server. The two primary benefits of a proxy server are to improve performance and to filter requests. If a proxy server is required, then the proxy server is installed 801. The process software is sent to the servers either via a protocol such as FTP or it is copied directly from the source files to the server files via file sharing 802. Another embodiment would be to send a transaction to the servers that contain the process software and have the server process the transaction, then receive and copy the process software to the server's file system. Once the process software is stored at the servers, the users, via their client computers, then access the process software on the servers and copy it to their client computers file systems 803. Another embodiment is to have the servers automatically copy the process software to each client and then run the installation program for the process software at each client computer. The users execute the program that installs the process software on their client computer 812, and then exit the process 708.

In step 704, a determination is made whether the process software is to be deployed by sending the process software to users via e-mail. The set of users where the process software will be deployed are identified together with the addresses of the user client computers 705. The process software is sent via e-mail 804 to each of the users' client computers. The users receive the e-mail 805 and then detach the process software from the e-mail to a directory on their client computers 806. The users execute the program that installs the process software on their client computer 812, and then exit the process 708.

Lastly, a determination is made on whether to the process software will be sent directly to user directories on their client computers 706. If so, the user directories are identified 707. The process software is transferred directly to the users' client computer directory 807. This can be done in several ways such as but not limited to sharing of the file system directories and then copying from the sender's file system to the recipient users' file system or alternatively using a transfer protocol such as File Transfer Protocol (FTP). The users access the directories on their client file systems in preparation for installing the process software 808. The users execute the program that installs the process software on their client computer 812, and then exit the process 708.

The process software is integrated into a client, server, and network environment by providing for the process software to coexist with applications, operating systems, and network operating systems software, and then installing the process software on the clients and servers in the environment where the process software will function.

The first step is to identify any software on the clients and servers including the network operating system where the process software will be deployed that are required by the process software or that work in conjunction with the process software. This includes the network operating system that is software that enhances a basic operating system by adding networking features.

Next, the software applications and version numbers will be identified and compared to the list of software applications and version numbers that have been tested to work with the process software. Those software applications that are missing or that do not match the correct version will be upgraded with the correct version numbers. Program instructions that pass parameters from the process software to the software applications will be checked to ensure the parameter lists match the parameter lists required by the process software. Conversely, parameters passed by the software applications to the process software will be checked to ensure the parameters match the parameters required by the process software. The client and server operating systems including the network operating systems will be identified and compared to the list of operating systems, version numbers, and network software that have been tested to work with the process software. Those operating systems, version numbers, and network software that do not match the list of tested operating systems and version numbers will be upgraded on the clients and servers to the required level.

After ensuring that the software, where the process software is to be deployed, is at the correct version level that has been tested to work with the process software, the integration is completed by installing the process software on the clients and servers.

In FIG. 13, step 820 begins the integration of the process software. The first thing is to determine if there are any process software programs that will execute on a server or servers 821. If this is not the case, then integration proceeds to 827. If this is the case, then the server addresses are identified

822. The servers are checked to see if they contain software that includes the operating system (OS), applications, and network operating systems (NOS), together with their version numbers, that have been tested with the process software **823**. The servers are also checked to determine if there is any missing software that is required by the process software **823**.

A determination is made if the version numbers match the version numbers of OS, applications, and NOS that have been tested with the process software **824**. If all of the versions match and there is no missing required software, the integration continues in **827**.

If one or more of the version numbers do not match, then the unmatched versions are updated on the server or servers with the correct versions **825**. Additionally, if there is missing required software, then it is updated on the server or servers **825**. The server integration is completed by installing the process software **826**.

Step **827**, which follows either step **821**, **824**, or **826**, determines if there are any programs of the process software that will execute on the clients. If no process software programs execute on the clients, the integration proceeds to **830** and exits. If this is not the case, then the client addresses are identified at **828**.

The clients are checked to see if they contain software that includes the operating system (OS), applications, and network operating systems (NOS), together with their version numbers, that have been tested with the process software **829**. The clients are also checked to determine if there is any missing software that is required by the process software **829**.

A determination is made as to whether the version numbers match the version numbers of OS, applications, and NOS that have been tested with the process software **831**. If all of the versions match and there is no missing required software, then the integration proceeds to **830** and exits.

If one or more of the version numbers do not match, then the unmatched versions are updated on the clients with the correct versions **832**. In addition, if there is missing required software then it is updated on the clients **832**. The client integration is completed by installing the process software on the clients **833**. The integration proceeds to **830** and exits.

The process software can be stored on a shared file system accessible from one or more servers. The process software is executed via transactions that contain data and server processing requests that use CPU units on the accessed server. CPU units are units of time such as minutes, seconds, hours on the central processor of the server. Additionally, the assessed server may make requests of other servers that require CPU units. CPU units are an example that represents one measurement of use. Other measurements of use include but are not limited to network bandwidth, memory usage, storage usage, packet transfers, complete transactions etc. When multiple customers use the same process software application, their transactions are differentiated by the parameters included in the transactions that identify the unique customer and the type of service for that customer. All of the CPU units and other measurements of use that are used for the services for each customer are recorded. When the number of transactions to any one server reaches a number that begins to affect the performance of that server, other servers are accessed to increase the capacity and to share the workload. Likewise, when other measurements of use such as network bandwidth, memory usage, storage usage, etc. approach a capacity so as to affect performance, additional network bandwidth, memory usage, storage etc. are added to share the workload. The measurements of use used for each service and customer are sent to a collecting server that sums the measurements of use for each customer for each service that was

processed anywhere in the network of servers that provide the shared execution of the process software. The summed measurements of use units are periodically multiplied by unit costs and the resulting total process software application service costs are alternatively sent to the customer and/or indicated on a web site accessed by the customer, which then remits payment to the service provider. In another embodiment, the service provider requests payment directly from a customer account at a banking or financial institution. In another embodiment, if the service provider is also a customer of the customer that uses the process software application, the payment owed to the service provider is reconciled to the payment owed by the service provider to minimize the transfer of payments.

The process software is shared, simultaneously serving multiple customers in a flexible, automated fashion. It is standardized, requiring little customization, and it is scalable, providing capacity on demand in a pay-as-you-go model.

The process software can be stored on a shared file system accessible from one or more servers. The process software is executed via transactions that contain data and server processing requests that use CPU units on the accessed server. CPU units are units of time such as minutes, seconds, hours on the central processor of the server. Additionally, the assessed server may make requests of other servers that require CPU units. CPU units are an example that represents one measurement of use. Other measurements of use include but are not limited to network bandwidth, memory usage, storage usage, packet transfers, complete transactions etc.

When multiple customers use the same process software application, their transactions are differentiated by the parameters included in the transactions that identify the unique customer and the type of service for that customer. All of the CPU units and other measurements of use that are used for the services for each customer are recorded. When the number of transactions to any one server reaches a number that begins to affect the performance of that server, other servers are accessed to increase the capacity and to share the workload. Likewise, when other measurements of use such as network bandwidth, memory usage, storage usage, etc. approach a capacity so as to affect performance, additional network bandwidth, memory usage, storage etc. are added to share the workload.

The measurements of use used for each service and customer are sent to a collecting server that sums the measurements of use for each customer for each service that was processed anywhere in the network of servers that provide the shared execution of the process software. The summed measurements of use units are periodically multiplied by unit costs and the resulting total process software application service costs are alternatively sent to the customer and/or indicated on a web site accessed by the customer, which then remits payment to the service provider.

In another embodiment, the service provider requests payment directly from a customer account at a banking or financial institution.

In another embodiment, if the service provider is also a customer of the customer that uses the process software application, the payment owed to the service provider is reconciled to the payment owed by the service provider to minimize the transfer of payments.

In FIG. 14, step **840** begins the On Demand process. A transaction is created that contains the unique customer identification, the requested service type, and any service parameters that further specify the type of service **841**. The transaction is then sent to the main server **842**. In an On Demand

environment, the main server can initially be the only server, then, as capacity is consumed, other servers are added to the On Demand environment.

The server central processing unit (CPU) capacities in the On Demand environment are queried **843**. The CPU requirement of the transaction is estimated, then the servers' available CPU capacity in the On Demand environment are compared to the transaction CPU requirement to see if there is sufficient CPU capacity available in any server to process the transaction **844**. If there is not sufficient server CPU capacity available, then additional server CPU capacity is allocated to process the transaction **848**. If there was already sufficient CPU capacity available, then the transaction is sent to a selected server **845**.

Before executing the transaction, a check is made of the remaining On Demand environment to determine if the environment has sufficient available capacity for processing the transaction. This environment capacity consists of such things as, but not limited to, network bandwidth, processor memory, storage etc. **846**. If there is not sufficient available capacity, then capacity will be added to the On Demand environment **847**. Next, the required software to process the transaction is accessed, loaded into memory, then the transaction is executed **849**.

The usage measurements are recorded **850**. The usage measurements consist of the portions of those functions in the On Demand environment that are used to process the transaction. The usage of such functions as, but not limited to, network bandwidth, processor memory, storage and CPU cycles are what is recorded. The usage measurements are summed, multiplied by unit costs, and then recorded as a charge to the requesting customer **851**. If the customer has requested that the On Demand costs be posted to a web site **852**, then they are posted **853**.

If the customer has requested that the On Demand costs be sent via e-mail to a customer address **854**, then they are sent **855**. If the customer has requested that the On Demand costs be paid directly from a customer account **856**, then payment is received directly from the customer account **857**. The last step is to exit the On Demand process **858**.

The process software may be deployed, accessed and executed through the use of a virtual private network (VPN), which is any combination of technologies that can be used to secure a connection through an otherwise unsecured or untrusted network. The use of VPNs is to improve security and for reduced operational costs. The VPN makes use of a public network, usually the Internet, to connect remote sites or users together. Instead of using a dedicated, real-world connection such as leased line, the VPN uses "virtual" connections routed through the Internet from the company's private network to the remote site or employee.

The process software may be deployed, accessed, and executed through either a remote-access or a site-to-site VPN. When using the remote-access VPNs, the process software is deployed, accessed, and executed via the secure, encrypted connections between a company's private network and remote users through a third-party service provider. The enterprise service provider (ESP) sets a network access server (NAS) and provides the remote users with desktop client software for their computers. The telecommuters can then dial a toll-free number or attach directly via a cable or DSL modem to reach the NAS and use their VPN client software to access the corporate network and to access, download, and execute the process software.

When using the site-to-site VPN, the process software is deployed, accessed, and executed through the use of dedi-

cated equipment and large-scale encryption, which are used to connect a company's multiple fixed sites over a public network, such as the Internet.

The process software is transported over the VPN via tunneling, which is the process of placing an entire packet within another packet and sending it over a network. The protocol of the outer packet is understood by the network and both points, called tunnel interfaces, where the packet enters and exits the network.

In FIGS. **15** and **16**, step **860** begins the Virtual Private Network (VPN) process. A determination is made to see if a VPN for remote access is required **861**. If it is not required, then proceed to **862**. If it is required, then determine if the remote access VPN exists **864**.

If it does exist, then proceed to **865**. Otherwise, identify the third party provider that will provide the secure, encrypted connections between the company's private network and the company's remote users **876**. The company's remote users are identified **877**. The third party provider then sets up a network access server (NAS) **878** that allows the remote users to dial a toll-free number or attach directly via a cable or DSL modem to access, download, and install the desktop client software for the remote-access VPN **879**.

After the remote access VPN has been built, or if it been previously installed, the remote users can then access the process software by dialing into the NAS or attaching directly via a cable or DSL modem into the NAS **865**. This allows entry into the corporate network where the process software is accessed **866**. The process software is transported to the remote users' desktop over the network via tunneling. That is, the process software is divided into packets and each packet, including the data and protocol, is placed within another packet **867**. When the process software arrives at the remote users' desktop, it is removed from the packets, reconstituted, and executed on the remote users' desktop **868**.

A determination is made to see if a VPN for site-to-site access is required **862**. If it is not required, then proceed to exit the process **863**. Otherwise, determine if the site-to-site VPN exists **869**. If it does exist, then proceed to **872**. Otherwise, install the dedicated equipment required to establish a site-to-site VPN **870**. Then build the large-scale encryption into the VPN **871**.

After the site-to-site VPN has been built, or if it had been previously established, the users access the process software via the VPN **872**. The process software is transported to the site users over the network via tunneling **873**. That is, the process software is divided into packets and each packet, including the data and protocol, is placed within another packet **874**. When the process software arrives at the remote users' desktop, it is removed from the packets, reconstituted, and executed on the site users' desktop **875**. Proceed to exit the process **863**.

The terminology used herein is for the purpose of describing particular embodiments only and is not intended to be limiting of this disclosure. As used herein, the singular forms "a", "an", and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

The corresponding structures, materials, acts, and equivalents of all means or step plus function elements in the claims below are intended to include any structure, material, or act for performing the function in combination with other

claimed elements as specifically claimed. The descriptions of the various embodiments herein have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

What is claimed is:

1. A method of determining that a user is human, said method comprising:

receiving a first question into a question-answer system; generating candidate answers to said first question and confidence scores associated with said candidate answers, said confidence scores comprising a measure of the likelihood a candidate answer is correct; identifying missing information that improves a confidence score for at least one of said candidate answers to said first question; generating a user question in natural language format, to provide said missing information, based on said first question; outputting said user question to a user; receiving a user response to said user question from said user; and validating said user response as having been provided by a human based on characteristics of said user response.

2. The method of claim 1, said validating said user response comprising comparing said user response with a response of a trusted source, said user response having been provided by a human if said user response matches said response of said trusted source.

3. The method of claim 1, said validating said user response comprising comparing said user response with a previously validated response, said user response having been provided by a human if said user response matches said previously validated response.

4. The method of claim 1, said validating said user response comprising comparing said user response with a plurality of responses from a community source separate from said question-answer system, said user response having been provided by a human if said user response matches a threshold number of responses from said community source.

5. The method of claim 1, said validating said user response comprising determining if said user response is in natural language format.

6. The method of claim 5, said determining if said user response is in natural language format comprising evaluating syntax of said user response.

7. The method of claim 6, said evaluating said syntax comprising parsing said user response into syntactic elements and determining if said user response conforms to syntactic rules of a language.

8. The method of claim 1, said missing information comprising one of a data, a fact, a syntactical relationship, a grammatical relationship, a logical rule, and a taxonomy rule.

9. The method of claim 1, said identifying missing information comprising:

evaluating a piece of evidence relevant to at least one of said candidate answers in relation to said first question in order to identify said missing information.

10. The method of claim 9, said evaluating a piece of evidence comprising:

parsing said first question into a collection of elements; parsing said piece of evidence into a collection of elements; analyzing said collection of elements for said first question and said piece of evidence in order to determine a relationship between an element of said collection of elements for said first question and an element of said collection of elements for said piece of evidence; and locating a missing relationship between said elements, said missing relationship comprising said missing information.

11. The method of claim 10, said relationship between elements comprising one of a lexical relationship, a grammatical relationship, and a semantic relationship.

12. A method of determining that a user of a question-answer system is human, said method comprising:

receiving a question into a computerized question-answer system operating on a computerized device to provide an automatically generated answer to said question; automatically generating, by said computerized device, a plurality of candidate answers to said question; evaluating, by said computerized device, sources of evidence used to generate said plurality of candidate answers to identify marginal evidence, said marginal evidence contributing only partially to a candidate answer;

determining, by said computerized device, a confidence score for each of said plurality of candidate answers, said confidence score comprising a measure of the likelihood a candidate answer is correct;

automatically identifying, by said computerized device, information not provided by said marginal evidence that could further develop said confidence score;

automatically generating, by said computerized device, at least one follow-on inquiry based on said information; outputting, by said computerized device, said at least one follow-on inquiry to external sources separate from said question-answer system to obtain responses to said follow-on inquiry;

receiving, by said computerized device, from said external sources at least one response to said at least one inquiry; and

automatically evaluating, by said computerized device, whether said at least one response to said at least one follow-on inquiry was provided by a human respondent based on one of:

a previous determination that a respondent is human, and syntactic or semantic characteristics of said at least one response.

13. The method of claim 12, further comprising: tracking answers and responses to previous questions or inquiries from external sources known to be humans producing tracked answers, using said computerized device;

comparing a response received from said respondent to said tracked answers and responses, using said computerized device; and

determining that said respondent is human when said response received from said respondent matches said tracked answers and responses, using said computerized device.

14. The method of claim 12, said generating at least one follow-on inquiry comprising forming a natural-language inquiry; and

said evaluating whether said at least one response to said at least one follow-on inquiry was provided by a human respondent comprising one of assessing accuracy of said

31

at least one response, determining whether or not said at least one response is in natural language format, and any combination of said assessing accuracy of said at least one response and said determining whether or not said at least one response is in natural language format.

15. The method of claim 12, said generating at least one follow-on inquiry comprising forming an inquiry that humans have a known particular manner of responding; and said evaluating whether said at least one response to said at least one follow-on inquiry was provided by a human respondent comprising determining that said respondent's response matches said known particular manner of responding.

16. The method of claim 12, said evaluating whether said at least one response was provided by a human respondent comprising finding a most common response obtained from said external sources to said inquiry.

17. A method comprising:

providing a first question to be answered by a Question Answering (QA) system to a processor;

creating, by said processor, a collection of candidate answers to said first question, said collection of candidate answers being created from a corpus of data;

generating, by said processor, supporting evidence and a confidence score for each of said candidate answers, said confidence score comprising a measure of the likelihood a candidate answer is correct;

identifying, by an evidence analysis module, information not provided by said supporting evidence that could further develop said confidence score;

producing, by said evidence analysis module, a second question based on said information;

presenting, by said processor, said second question to external sources separate from said QA system to obtain responses to said second question;

receiving, by said processor, at least one response or knowledge item from said external sources; and

determining, by said processor, whether a respondent is human based on one of:

a previous determination that said respondent is human, and

syntactic or semantic characteristics of said at least one response or knowledge item provided to said second question.

18. The method of claim 17, further comprising:

tracking, by said processor, answers to previous questions received from external sources previously known to be humans, producing tracked answers;

comparing, by said processor, said at least one response received from a respondent to said tracked answers; and

determining, by said processor, that said respondent is human when said at least one response received from said respondent matches said tracked answers.

19. The method of claim 17, said producing a second question comprising forming a natural-language question; and

said determining whether said respondent is human comprising one of assessing accuracy of said at least one response, determining whether or not said at least one response is in natural language format, and any combination of said assessing accuracy of said at least one response and said determining whether or not said at least one response is in natural language format.

20. A method comprising:

receiving a question into a Question Answering (QA) system;

32

said QA system comparing said question to a corpus of data;

said QA system generating hypotheses about relationships between linguistic and semantic entities of said question and said corpus of data;

said QA system generating candidate answers to said question using said hypotheses;

said QA system evaluating sources of evidence used to generate said candidate answers to identify marginal evidence, said marginal evidence contributing only partially to a candidate answer;

said QA system determining a confidence score for each of said candidate answers, said confidence score comprising a measure of the likelihood a candidate answer is correct;

said QA system determining a missing piece of knowledge comprising information that would enable said QA system to further develop said confidence score;

said QA system formulating a follow-on inquiry to obtain said missing piece of knowledge;

said QA system outputting said follow-on inquiry to an external community source separate from said QA system to obtain responses to said follow-on inquiry;

said QA system receiving from said external community source responses to said follow-on inquiry comprising said missing piece of knowledge;

said QA system generating a knowledge inquiry;

said QA system outputting said knowledge inquiry to said external community source;

said QA system receiving from said external community source a response to said knowledge inquiry; and

said QA system determining whether a respondent is human based on syntactic or semantic characteristics of said response to said knowledge inquiry.

21. The method of claim 20, said knowledge inquiry requiring one of semantic and linguistic knowledge to obtain a correct response.

22. The method of claim 20, said determining whether a respondent is human comprising comparing said response to said knowledge inquiry with other responses obtained for said knowledge inquiry and finding a most common response among various respondents regarding responses to said knowledge inquiry.

23. The method of claim 20, said generating a knowledge inquiry comprising forming a natural-language inquiry; and said determining whether said respondent is human comprising one of assessing accuracy of said response to said knowledge inquiry, determining whether or not said response to said knowledge inquiry is in natural language format, and any combination of said assessing accuracy of said response and said determining whether or not said response is in natural language format.

24. The method of claim 20, said method further comprising:

said QA system extracting evidence related to said question from said corpus of data; and

said QA system classifying said evidence as one of good evidence and marginal evidence,

said good evidence having an evidence score above a previously established evidence threshold value, and enabling said QA system to provide a candidate answer to said question with a confidence score above a previously established confidence threshold value, and

said marginal evidence having an evidence score below said previously established evidence threshold value, enabling said QA system to provide a candidate answer to said question with a confidence score below said

previously established confidence threshold value, and requiring said QA system to obtain additional information to improve said confidence score.

* * * * *