



US 20180137433A1

(19) **United States**(12) **Patent Application Publication****Devarakonda et al.**(10) **Pub. No.: US 2018/0137433 A1**(43) **Pub. Date: May 17, 2018**(54) **SELF-TRAINING OF QUESTION
ANSWERING SYSTEM USING QUESTION
PROFILES**(71) Applicant: **International Business Machines
Corporation**, Armonk, NY (US)(72) Inventors: **Murthy V. Devarakonda**, Peekskill,
NY (US); **Siddharth A. Patwardhan**,
Yorktown Heights, NY (US); **Preethi
Raghavan**, Cambridge, MA (US)(21) Appl. No.: **15/353,172**(22) Filed: **Nov. 16, 2016****Publication Classification**(51) **Int. Cl.**
G06N 99/00 (2006.01)
G09B 5/12 (2006.01)
G06N 5/02 (2006.01)
G06F 19/00 (2006.01)(52) **U.S. Cl.**
CPC **G06N 99/005** (2013.01); **G06F 19/322**
(2013.01); **G06N 5/022** (2013.01); **G09B 5/12**
(2013.01)(57) **ABSTRACT**

Mechanisms are provided to implement a self-training engine of a question and answer system. The self-training engine pairs an unanswered natural language question with portions of an electronic document to generate an unlabeled data set. The self-training engine trains a model based on a labeled data set comprising a finite number of question and answer pair data structures and applies the model to the unlabeled data set to identify an answer from the portions of the electronic document to the unanswered natural language question. The self-training engine modifies the labeled data set to include the identified answer and corresponding unanswered natural language question as an additional question and answer pair data structure. The self-training engine then trains the model based on the modified labeled data set.




FIG. 1

QUESTION: WHEN WAS THE PATIENT'S LAST LIPID PANEL PERFORMED?

210

QUESTION PROFILE

QUESTION TYPE: FACT-FINDING
QUESTION TEMPLATE: WHEN WAS <TEST> DONE?

ANSWER TYPE: TEMPORAL
KEY ENTITY: LIPID PANEL
KEY ENTITY TYPE: TEST
KEY ENTITY MODIFIER: LAST
NUMBER OF ANSWERS: SINGLE
ANSWER NOTE TYPES: LAB REPORTS, STRUCTURED PROCEDURES

220

EVENLY DISTRIBUTED ACROSS ALL RECENT LAB NOTES
CARDINALITY: 1

FIG. 2A

QUESTION: WHY WERE X-RAYS DONE ON THE FOREARM AND HAND?

230

QUESTION PROFILE

QUESTION TYPE: REASON
QUESTION TEMPLATE: WHY WAS <TEST> DONE ON <BODY_PARTS>?

ANSWER TYPE: PASSAGE, NOTES
KEY ENTITY: X-RAYS
KEY ENTITY TYPE: TEST
NUMBER OF ANSWERS: MULTIPLE
ANSWER NOTE TYPES: PROGRESS NOTES, DISCHARGE SUMMARIES, LAB REPORTS

240

EVENLY DISTRIBUTED ACROSS PROGRESS NOTES AND DISCHARGE SUMMARIES AROUND THE TIME OF THE LAB REPORT THAT MENTIONS THIS X-RAY
CARDINALITY: 5

FIG. 2B




FIG. 3




FIG. 4





FIG. 5



SELF-TRAINING OF QUESTION ANSWERING SYSTEM USING QUESTION PROFILES

BACKGROUND

[0001] The present application relates generally to an improved data processing apparatus and method and more specifically to mechanisms for providing a self-training functionality for a question answering system using question profiles.

[0002] Decision-support systems exist in many different industries where human experts require assistance in retrieving and analyzing information. An example that will be used throughout this application is a diagnosis system employed in the healthcare industry. Diagnosis systems can be classified into systems that use structured knowledge, systems that use unstructured knowledge, and systems that use clinical decision formulas, rules, trees, or algorithms. The earliest diagnosis systems used structured knowledge or classical, manually constructed knowledge bases. The Internist-I system developed in the 1970s uses disease-finding relations and disease-disease relations. The MYCIN system for diagnosing infectious diseases, also developed in the 1970s, uses structured knowledge in the form of production rules, stating that if certain facts are true, then one can conclude certain other facts with a given certainty factor. DXplain, developed starting in the 1980s, uses structured knowledge similar to that of Internist-I, but adds a hierarchical lexicon of findings.

[0003] Iliad, developed starting in the 1990s, adds more sophisticated probabilistic reasoning where each disease has an associated a priori probability of the disease (in the population for which Iliad was designed), and a list of findings along with the fraction of patients with the disease who have the finding (sensitivity), and the fraction of patients without the disease who have the finding (1-specificity).

[0004] In 2000, diagnosis systems using unstructured knowledge started to appear. These systems use some structuring of knowledge such as, for example, entities such as findings and disorders being tagged in documents to facilitate retrieval. ISABEL, for example, uses Autonomy information retrieval software and a database of medical textbooks to retrieve appropriate diagnoses given input findings. Autonomy Aumience uses the Autonomy technology to retrieve diagnoses given findings and organizes the diagnoses by body system. First CONSULT allows one to search a large collection of medical books, journals, and guidelines by chief complaints and age group to arrive at possible diagnoses. PEPID DDX is a diagnosis generator based on PEPID's independent clinical content.

[0005] Clinical decision rules have been developed for a number of medical disorders, and computer systems have been developed to help practitioners and patients apply these rules. The Acute Cardiac Ischemia Time-Insensitive Predictive Instrument (ACI-TIPI) takes clinical and ECG features as input and produces probability of acute cardiac ischemia as output to assist with triage of patients with chest pain or other symptoms suggestive of acute cardiac ischemia. ACI-TIPI is incorporated into many commercial heart monitors/defibrillators. The CaseWalker system uses a four-item questionnaire to diagnose major depressive disorder. The PKC Advisor provides guidance on 98 patient problems such as abdominal pain and vomiting.

SUMMARY

[0006] This Summary is provided to introduce a selection of concepts in a simplified form that are further described herein in the Detailed Description. This Summary is not intended to identify key factors or essential features of the claimed subject matter, nor is it intended to be used to limit the scope of the claimed subject matter.

[0007] In one illustrative embodiment, a method is provided, in a data processing system comprising a processor and a memory, the memory comprising instructions which cause the processor to be configured to implement a self-training engine of a question and answer system. The method comprises pairing, by the self-training engine, an unanswered natural language question with portions of an electronic document to generate an unlabeled data set. The method further comprises training, by the self-training engine, a model based on a labeled data set comprising a finite number of question and answer pair data structures. Moreover, the method comprises applying, by the self-training engine, the model to the unlabeled data set to identify an answer from the portions of the electronic document to the unanswered natural language question. In addition, the method comprises modifying, by the self-training engine, the labeled data set to include the identified answer and corresponding unanswered natural language question as an additional question and answer pair data structure. Furthermore, the method comprises training, by the self-training engine, the model based on the modified labeled data set.

[0008] In other illustrative embodiments, a computer program product comprising a computer useable or readable medium having a computer readable program is provided. The computer readable program, when executed on a computing device, causes the computing device to perform various ones of, and combinations of, the operations outlined above with regard to the method illustrative embodiment.

[0009] In yet another illustrative embodiment, a system/apparatus is provided. The system/apparatus may comprise one or more processors and a memory coupled to the one or more processors. The memory may comprise instructions which, when executed by the one or more processors, cause the one or more processors to perform various ones of, and combinations of, the operations outlined above with regard to the method illustrative embodiment.

[0010] These and other features and advantages of the present invention will be described in, or will become apparent to those of ordinary skill in the art in view of, the following detailed description of the example embodiments of the present invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0011] The invention, as well as a preferred mode of use and further objectives and advantages thereof, will best be understood by reference to the following detailed description of illustrative embodiments when read in conjunction with the accompanying drawings, wherein:

[0012] FIG. 1 depicts an overview of this self-training functionality using question profiles in accordance with one illustrative embodiment;

[0013] FIGS. 2A and 2B illustrate examples of question profiles in accordance with one illustrative embodiment;

[0014] FIG. 3 depicts a schematic diagram of one illustrative embodiment of a cognitive system in a computer network, in which a question profile based self-training engine is implemented to train one or more models utilized by the cognitive system;

[0015] FIG. 4 is a block diagram of an example data processing system in which aspects of the illustrative embodiments are implemented;

[0016] FIG. 5 is an example diagram illustrating an interaction of elements of a healthcare cognitive system in accordance with one illustrative embodiment; and

[0017] FIG. 6 is a flowchart outlining an example operation for performing an iterative question profile based self-training operation in accordance with one illustrative embodiment.

DETAILED DESCRIPTION

[0018] The strengths of current cognitive systems, such as current medical diagnosis, patient health management, patient treatment recommendation systems, law enforcement investigation systems, and other decision support systems, are that they can provide insights that improve the decision making performed by human beings. For example, in the medical context, such cognitive systems may improve medical practitioners' diagnostic hypotheses, can help medical practitioners avoid missing important diagnoses, and can assist medical practitioners with determining appropriate treatments for specific diseases. However, current systems still suffer from significant drawbacks which should be addressed in order to make such systems more accurate and usable for a variety of applications as well as more representative of the way in which human beings make decisions, such as diagnosing and treating patients. In particular, one drawback of current systems is the considerable amount of human effort and resource expenditure needed to train such cognitive systems to perform their cognitive operations.

[0019] Cognitive systems, such as natural language based Question Answering (QA) based cognitive systems, are trained using machine learning technology that requires text that is annotated by human subject matter experts (SMEs). In order to successfully train a machine-learning model of a cognitive system, e.g., QA based cognitive system, the training set, e.g., a set of training questions and corresponding known (ground-truth) answers, used to perform the training need to be realistic and representative of requests/questions that would be submitted/asked by a user of the cognitive system. For purposes of the following discussion, it will be assumed that the cognitive system is a QA system but it should be appreciated that any cognitive system that processes requests/questions may be used without departing from the spirit and scope of the illustrative embodiments. The set of training questions and the corresponding answers are used to train the QA based cognitive system by modifying operational parameters of the QA based cognitive system so that it generates the correct given answer for the training question.

[0020] In the medical domain, for a patient-specific QA based cognitive system (hereafter referred to as a QA system), the requirement to have questions that are representative of actual questions that would be asked by the patient, or the medical professional treating the patient, makes it especially challenging and expensive to obtain a set of such training question-answer (QA) pairs. While it is challenging to automatically gather questions that a medical

professional, e.g., the patient's physician, would ask about patients during the course of medical professional's day, manually identifying correct answers to these questions (words, phrases or passages in the patient's medical record) is an even more time-consuming task and resource intensive task.

[0021] Thus, overall, the training of a QA system takes considerable human effort and time, especially in the medical domain, although similar levels of human effort and time may be present in other domains as well. In order to ensure that all aspects of a patient's electronic medical records (EMRs) are covered by the QA system, a large set of training questions and ground truth answers (typically tens of thousands of question-answer (QA) pairs) must be generated and used to train the QA system. This takes considerable time and effort to accomplish. Through studies it has been noted that the process of ground truth answer generation for training a medical domain QA system from a patient's EMRs may take approximately a year of human effort to generate ground truth answers for approximately 1750 training questions with a team of 11 SMEs (medical expert annotators) working on the task intermittently. In comparison, the question generation aspect of generating QA pairs for training the QA system is relatively faster with a single SME being able to generate approximately 15 questions per hour. This results in a relatively small training set of QA pairs and a larger set of unanswered questions. The challenge is to expand the relatively small labeled training set of QA pairs effectively without having to rely on a human based annotation effort. It should be appreciated that a "labeled" training set refers to the QA pair having an associated label indicating whether or not the "answer" of the QA pair answers the corresponding "question" of the QA pair. The label is often a binary 1 or 0 value where 1 indicates that the "answer" answers the "question" and a 0 indicates that the "answer" does not answer the "question" in the QA pair.

[0022] Self-training is a method that has been successfully used to address limited training data on many natural language processing (NLP) tasks, such as parsing, information extraction, word sense disambiguation, and the like. However, the self-training methodology has not been applied directly to the QA pair generation task for self-training a QA system, and especially not with regard to the many challenges specific to the medical domain and handling questions directed to, or answerable by, patient EMRs. Even in other domains, self-training has not been applied to the QA pair generation problem in the manner set forth hereafter with regard to the illustrative embodiments described herein.

[0023] The illustrative embodiments provide mechanisms for performing self-training of a QA system by pairing unanswered questions with portions of documentation to generate new QA pairs used for training the QA system. In the following description of the illustrative embodiments, various terms will be utilized having the following definitions specific to the illustrative embodiments:

1. Labeled Set (or Labeled Training Set): Questions that have been associated with correct answers. An example <question, answer> pair in the labeled training set is as follows:

<What medications does the patient take for diabetes?, a regular dose of insulin> The label corresponding to this

instance would be 1, thereby indicating that the answer correctly answers the question.

2. Unlabeled Set: Questions that have not yet been associated with correct answers and thus, do not have an associated label indicating that an answer is correct or not correct for the question.

3. Self-Training: A semi-supervised learning method for utilizing unlabeled data with labeled data to create a better machine learning model.

4. Question Profile: Attributes of the question and expected answer.

5. QA Pairs: A natural language question paired with a correct answer.

6. Answer: A fact, one or more passages of natural language text, one or more documents of natural language text, which is responsive to a natural language question.

7. QA Data Set/QA Training Set: A set of natural language questions along with their correct answers.

8. Corpus/Corpora: A collection of text documents covering information from a wide variety of topics and sources, e.g., a clinical corpus may contain electronic medical records for different patients potentially generated by, or otherwise available from, various sources.

9. Electronic Medical Records (EMRs): A collection of medical information for a patient stored in electronic form as one or more data structures. The medical information may be obtained from one or more sources of information and may represent various types of information generally directed to the medical condition of the patient at one or more points in time.

10. Clinical Notes: Portions of an EMR that represent results, observations, diagnosis, treatment instructions, or the like, generated by one or more medical personnel, medical equipment, or the like, and added to the EMR as a result of an interaction with a patient for the purposes of treating the patient. Clinical notes may be provided as natural language textual notes added to the EMR or may be provided in a more structured format. In some cases, clinical notes may be added to the patient's EMRs automatically by automated medical equipment.

[0024] With these definitions in mind, in the context of a medical domain based QA system, the illustrative embodiments generate QA pairs and perform self-training of a QA system using such QA pairs by pairing unanswered questions with clinical notes from a patient's electronic medical record (EMR). It should be appreciated that while medical domain based QA system self-training will be used as a basis for describing the operation of the illustrative embodiments, the present invention is not limited to this domain or such embodiments. To the contrary, the mechanisms of the illustrative embodiments may be utilized in any domain depending on the desired implementation, e.g., legal domain, financial domain, law enforcement domain, or any other domain where self-training of a QA system is desirable to avoid extensive expenditure of human resources and time to train the QA system.

[0025] Again, assuming a medical domain based QA system implementation, unanswered questions are paired with all of the clinical notes of the corresponding patient's EMR to form an unlabeled set of questions and answers. The self-training process trains a model on the limited labeled data, applies the semi-trained model on the unlabeled set, and selects one or more question-answer (QA) pairs for which corresponding confidence scores are sufficiently high,

e.g., equal to or greater than a predetermined threshold. These QA pairs are then added to the labeled set of QA pairs, where the labeled set of QA pairs is the training set used to train the model of the QA system. This process is iteratively repeated with incrementally expanding the labeled set of QA pairs until convergence, e.g., no more QA pairs having sufficiently high confidence are identified.

[0026] This iterative process for generating the training set of QA pairs automatically based on a limited size labeled set, to thereby iteratively expand the training set, and then utilizing the expanded training set of QA pairs to train the QA system, improves the QA system performance by providing a larger training set for training the QA system which captures more of the possible questions and corresponding answers that actual patients/medical professionals would ask the QA system during runtime and tuning the operation of the QA system to accommodate such. In addition, the illustrative embodiments further improve upon this process by incorporating question profiles into the process of selecting the QA pairs from the unlabeled set. The question profiles provide a mechanism for selecting high confidence scoring QA pairs (where the score is greater than a specific threshold) from the unlabeled set of QA pairs to be included in the labeled set of QA pairs. The profile for a question, or the "question profile," comprises the characteristics of the question and correct answer to the question, including, for example, key entities, relations between key entities, answer type, whether the answer is likely to occur across multiple portions of documents, e.g., notes in a patient's EMR, or a single portion, e.g., a single note in the patient EMR, and other pertinent characteristics. The question profile also stores information on all other questions in the question set (labeled and unlabeled questions) that are "similar" to this question, where similarity implies that both questions are essentially asking for the same answer, which may be determined from features extracted from the questions including focus, lexical answer type, and the like. Moreover, additional features indicative of similarity may also be utilized with logic specifically configured for identifying such instances of features and marking them as similarities between questions, e.g., temporally similar questions that map to identical time periods, questions that typically map to the same clinical note type around the same time period, questions where the answer is found in all clinical notes of the same clinical note type or same clinical note section, etc.

[0027] All of this information from a question's profile can be used to estimate the following: a) types of clinical notes where the answer is likely to be found, b) number of answer clinical notes corresponding to a question, and c) clusters of "similar" questions. This information can be used to influence the self-training process by setting a dynamic threshold based on the question's profile for selecting QA pairs from the unlabeled set, e.g., the dynamic threshold may be based on the type of clinical notes where the answer is likely to be found, the number of answer clinical notes, and/or thresholds associated with other similar questions in a cluster. This allows for considerable flexibility by allowing multiple clinical notes to be considered as an answer for certain question types. Further, these mechanisms can also be used for "label propagation", i.e. propagating labels or values indicating that an answer to a question actually answers the question or not, among questions that are from the same cluster when applying the trained model to the QA pairs in the unlabeled set.

[0028] Thus, the utilization of question profiles takes into account the longitudinal nature of the EMR data by allowing the system to choose a dynamic threshold and thus, finding the answer across multiple clinical notes. The implementation of question profiles also helps improve the overall accuracy of the system using label propagation. As noted above, the use of such question profiles is not limited to QA systems operating in the medical domain. The mechanisms of the illustrative embodiments may be implemented and applied to QA over any longitudinal dataset i.e. a dataset with related documents generated over time, for e.g., biographical articles, law documents, reports about evolving news events, financial documentation, law enforcement documentation, and/or the like.

[0029] Before beginning the discussion of the various aspects of the illustrative embodiments in more detail, it should first be appreciated that throughout this description the term “mechanism” will be used to refer to elements of the present invention that perform various operations, functions, and the like. A “mechanism,” as the term is used herein, may be an implementation of the functions or aspects of the illustrative embodiments in the form of an apparatus, a procedure, or a computer program product. In the case of a procedure, the procedure is implemented by one or more devices, apparatus, computers, data processing systems, or the like. In the case of a computer program product, the logic represented by computer code or instructions embodied in or on the computer program product is executed by one or more hardware devices in order to implement the functionality or perform the operations associated with the specific “mechanism.” Thus, the mechanisms described herein may be implemented as specialized hardware, software executing on general purpose hardware, software instructions stored on a medium such that the instructions are readily executable by specialized or general purpose hardware, a procedure or method for executing the functions, or a combination of any of the above.

[0030] The present description and claims may make use of the terms “a”, “at least one of”, and “one or more of” with regard to particular features and elements of the illustrative embodiments. It should be appreciated that these terms and phrases are intended to state that there is at least one of the particular feature or element present in the particular illustrative embodiment, but that more than one can also be present. That is, these terms/phrases are not intended to limit the description or claims to a single feature/element being present or require that a plurality of such features/elements be present. To the contrary, these terms/phrases only require at least a single feature/element with the possibility of a plurality of such features/elements being within the scope of the description and claims.

[0031] Moreover, it should be appreciated that the use of the term “engine,” if used herein with regard to describing embodiments and features of the invention, is not intended to be limiting of any particular implementation for accomplishing and/or performing the actions, steps, processes, etc., attributable to and/or performed by the engine. An engine may be, but is not limited to, software, hardware and/or firmware or any combination thereof that performs the specified functions including, but not limited to, any use of a general and/or specialized processor in combination with appropriate software loaded or stored in a machine readable memory and executed by the processor. Further, any name associated with a particular engine is, unless otherwise

specified, for purposes of convenience of reference and not intended to be limiting to a specific implementation. Additionally, any functionality attributed to an engine may be equally performed by multiple engines, incorporated into and/or combined with the functionality of another engine of the same or different type, or distributed across one or more engines of various configurations.

[0032] In addition, it should be appreciated that the following description uses a plurality of various examples for various elements of the illustrative embodiments to further illustrate example implementations of the illustrative embodiments and to aid in the understanding of the mechanisms of the illustrative embodiments. These examples intended to be non-limiting and are not exhaustive of the various possibilities for implementing the mechanisms of the illustrative embodiments. It will be apparent to those of ordinary skill in the art in view of the present description that there are many other alternative implementations for these various elements that may be utilized in addition to, or in replacement of, the examples provided herein without departing from the spirit and scope of the present invention.

[0033] The present invention may be a system, a method, and/or a computer program product. The computer program product may include a computer readable storage medium (or media) having computer readable program instructions thereon for causing a processor to carry out aspects of the present invention.

[0034] The computer readable storage medium can be a tangible device that can retain and store instructions for use by an instruction execution device. The computer readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable combination of the foregoing. A non-exhaustive list of more specific examples of the computer readable storage medium includes the following: 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), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing. A computer readable storage medium, as used herein, is not to be construed as being transitory signals per se, such as radio waves or other freely propagating electromagnetic waves, electromagnetic waves propagating through a waveguide or other transmission media (e.g., light pulses passing through a fiber-optic cable), or electrical signals transmitted through a wire.

[0035] Computer readable program instructions described herein can be downloaded to respective computing/processing devices from a computer readable storage medium or to an external computer or external storage device via a network, for example, the Internet, a local area network, a wide area network and/or a wireless network. The network may comprise copper transmission cables, optical transmission fibers, wireless transmission, routers, firewalls, switches, gateway computers and/or edge servers. A network adapter card or network interface in each computing/processing device receives computer readable program instructions from the network and forwards the computer readable

program instructions for storage in a computer readable storage medium within the respective computing/processing device.

[0036] Computer readable program instructions for carrying out operations of the present invention may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or either source code or object code 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 computer readable program instructions 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). In some embodiments, electronic circuitry including, for example, programmable logic circuitry, field-programmable gate arrays (FPGA), or programmable logic arrays (PLA) may execute the computer readable program instructions by utilizing state information of the computer readable program instructions to personalize the electronic circuitry, in order to perform aspects of the present invention.

[0037] Aspects of the present invention are described herein with reference to flowchart illustrations and/or block diagrams of methods, apparatus (systems), and computer program products according to embodiments of the invention. It will be understood that each block of the flowchart illustrations and/or block diagrams, and combinations of blocks in the flowchart illustrations and/or block diagrams, can be implemented by computer readable program instructions.

[0038] These computer readable 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. These computer readable program instructions may also be stored in a computer readable storage medium that can direct a computer, a programmable data processing apparatus, and/or other devices to function in a particular manner, such that the computer readable storage medium having instructions stored therein comprises an article of manufacture including instructions which implement aspects of the function/act specified in the flowchart and/or block diagram block or blocks.

[0039] The computer readable program instructions may also be loaded onto a computer, other programmable data processing apparatus, or other device to cause a series of operational steps to be performed on the computer, other programmable apparatus or other device to produce a computer implemented process, such that the instructions which execute on the computer, other programmable apparatus, or

other device implement the functions/acts specified in the flowchart and/or block diagram block or blocks.

[0040] The flowchart and block diagrams in the Figures illustrate the architecture, functionality, and operation of possible implementations of systems, methods, and computer program products according to various embodiments of the present invention. In this regard, each block in the flowchart or block diagrams may represent a module, segment, or portion of instructions, which comprises one or more executable instructions for implementing the specified logical function(s). 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 carry out combinations of special purpose hardware and computer instructions.

[0041] As noted above, the present invention provides mechanisms for implementing a self-training functionality for a question answering system using question profiles. An overview of this self-training functionality using question profiles, in accordance with one illustrative embodiment, is depicted in FIG. 1. As shown in FIG. 1, the training set size is iteratively increased by automatically assigning labels to unlabeled set (U) of question and answer pair candidates **120** using a machine learning (ML) model **140** trained in a previous iteration of the self-training process. The question profile **130** associated with each unanswered question in the unlabeled set U **120**, as determined from an association of the unanswered question to an answered question in the labeled training set (L) **110**, is used to influence the threshold and the number of question-answer (e.g., clinical note) pairs, through process **150**, that are added from the unlabeled set U **120** to the labeled training set (L) **110**, which is then used to train the machine learning model **140** for the next iteration.

[0042] Thus, given a small set of labeled QA pairs (L) **110** that can be considered as labeled training set for the machine learning (ML) model, a larger set of unanswered questions paired with each clinical note in the corresponding EMR, i.e. unlabeled set (U) **120**, and question profiles **130** associated with each question in the labeled training set, the detailed set of operations of the self-training process shown in FIG. 1 are as follows:

1. Train the machine learning model **140**: A model **140** is created using the labeled training set L **110**.
2. Apply the trained model **140** to the unlabeled set U **120**: the model **140** is applied to the unlabeled set U **120** of unanswered questions that are paired with all clinical notes of a patient EMR to generate confidence scores associated with each of these pairings, where the confidence scores indicate a confidence of the model **140** that the answer of a QA pair is a correct answer for the unanswered question of the QA pair. This generates, for each unanswered question, a set of answers and confidence scores associated with each of the answers, where the answers are clinical notes in the EMR.
3. For each question in the unlabeled set U **120**, identify a question profile: For each question in the unlabeled set U

120, a corresponding question profile is identified from the question profile set **130** that corresponds to the labeled training set **L 110**. That is, a similarity evaluation is performed for determining a similarity between the unanswered question in the unlabeled set **U 120** and a question in the labeled training set **110**. A similarity score based on a correlation of features may be generated and a highest scoring similar question in the labeled training set **L 110** is identified for the unanswered question in the unlabeled set **U 120**. The selected similar question in the labeled training set **L 110** has a corresponding question profile which is then applied to the unanswered question in the unlabeled set **U 120** to identify a corresponding confidence score threshold value and/or a number of QA pairs to select for inclusion in the labeled training set **L 110** for the unanswered question in the unlabeled set **U 120**.

4. For each question in the unlabeled set **U 120**, a dynamic threshold for selecting QA pairs based on their confidence scores is determined: In general, QA pairs are selected for each question based on the confidence score meeting or exceeding a threshold confidence score and thus, indicating a high confidence. Here, “high” confidence is determined by a threshold value, where typically, QA pairs meeting or exceeding a threshold confidence score value are considered to have a “high” confidence and can be added to the labeled training set **L 110**. The threshold value determines the confidence of the model **140** for a QA pair in the unlabeled set **U 120** being classified as positive, i.e. that the answer is a correct answer for the question. With the illustrative embodiments, instead of simply using a “fixed” threshold, a “dynamic” threshold may be used based on a question’s corresponding question profile.

[0043] In order to determine the threshold, information from the question profile for the question is used to determine what threshold confidence score value to utilize and/or the number of QA pairs to select. Characteristics of the question and its answer may be identified in the profile that indicate various criteria for determining a confidence score threshold/number of QA pairs to select. For example, if a question profile indicates that the answer is expected to be “evenly distributed across the EMR, with a cardinality of approximately 20 clinical notes”, this information may indicate that the top 20 QA pairs from the unlabeled set **U 120** to the labeled training set **L 110**. Similarly, if the question profile indicates that a particular question type typically has only 1 clinical note based answer associated with it, only the top scoring QA pair from the unlabeled set **U 120** is added to the labeled training set **L 110**. If the question profile for the question lacks this information, or if a sufficiently similar question is not found such that a corresponding question profile is not associated with the question, then a default operation is implemented that utilizes a fixed threshold confidence score value and/or corresponding default number of QA pairs are added from the unlabeled set **U 120** to the labeled training set **L 110**, where the value is tuned on a separate development set (small set of labeled data used for parameter tuning) **160**. In machine learning, a development set is a labeled dataset that is independent of the training set. This development set is used to tune parameters for the learning processing including threshold values.

5. Select QA pairs to add to labeled training set **L 110**: Based on the identified a dynamic threshold or default threshold, and the confidence scores associated with the QA pairs in the

unlabeled set **U 120** for the question, QA pairs that have confidence scores meeting or exceeding the threshold confidence score value are selected, via process **150**. The number of such QA pairs that are selected may be based on the dynamic determination corresponding to the question profile.

6. Label propagation: Labels are propagated to other questions belonging to a same cluster as indicated by characteristics in the question profile **130** associated with the question. That is, a clustering operation may be performed on questions in the unlabeled set **U 120** based on characteristics in the question profiles associated with those questions, or based on features extracted from those questions. For instance, a question cluster may contain a set of questions with a high semantic similarity score, i.e. all the questions in the same cluster are essentially asking for the same answer in different ways. One such question may get classified with a particular answer, e.g., clinical note, with high confidence as positive, e.g., a value of 1 indicating the answer is correct for the question. This positive label of the answer may now be assigned to all other questions in the question cluster with regard to this answer, thus increasing label accuracy of the QA pairs added from the unlabeled set **U 120** to the labeled training set **L 110**.

7. Update Trained Model based on Expanded Labeled Training Set **L' 110**: The trained machine learning model **140** is updated based on the newly expanded labeled training set **L'** which includes more QA pairs than a previous iteration.

[0044] The above process of operations 1-7 may be repeated until all the unanswered questions in the unlabeled set **U 120** are answered and corresponding QA pairs are added to the labeled training set **L 110**, or the training machine learning model **140** converges, i.e. the changes to parameters of operation of the model **140** do not change above a predetermined threshold value.

[0045] As noted above, the illustrative embodiments provide a self-training functionality that is based, at least in part, on question profiles **130** associated with questions in the labeled training set **L 110**. The question profiles **130** may be generated manually by a subject matter expert (SME). The question profiles, or at least portions of the question profiles, may be generated automatically using automated tools that may implement a machine learning algorithm that analyzes questions and corresponding answers in the labeled training set **L 110** to generate features and predict profile characteristics for other questions. Moreover, a rule-based algorithm may be implemented to analyze questions and predict profile characteristics, where the rules may be generated with the assistance of manual intervention by a SME, for example. Any automated, manual, or semi-automated methodology for generating questions profiles **130** may be used without departing from the spirit and scope of the present invention.

[0046] FIGS. 2A and 2B illustrate examples of question profiles in accordance with one illustrative embodiment. The question profile **220** in FIG. 2A is for a question **210** in the labeled training set **L 110** of the type “When was the patient’s last lipid panel performed?” The question profile **240** in FIG. 2B is for a question **230** in the labeled training set **L 110** of the type “Why were x-rays done on the forearm and hand?” The question profiles may be correlated with other similar questions, such as unanswered questions in the unlabeled set **U 120**, based on similarity measures between characteristics specified in the question profile with features extracted from the unanswered questions, or features

extracted from the answered questions **210** and **230** with features extracted from the unanswered questions.

[0047] As shown in FIGS. **2A** and **2B**, the question profile comprises a question type which indicates the goal of the question, e.g., to obtain a particular fact (fact-finding type), to determine the reason behind a fact (reason type), or the like. The question profile further comprises a question template which indicates the general format of the question with tags, e.g., “<test>”, “<body_parts>”, etc., indicating portions of the template that correspond to key entities. The question profile further comprises an answer type which indicates a classification of the nature of the answer, e.g., “temporal” meaning the answer is a date, time, or other temporal term/phrase, “passage” meaning that the answer is a passage from the EMR, “notes” meaning that the answer is found in clinical notes in the EMR, etc.

[0048] The question profile may also comprise one or more key entities and corresponding key entity types and modifiers which indicate the key features extracted from the question and characterize what the question is asking for, i.e. the key features that are indicative of the answer. For example, in FIG. **2A**, the key entity is “lipid panel” and the key entity type is “test,” with a key entity modifier of “last”, which is indicative of the answer being a temporal answer for the last lipid panel test. In FIG. **2B**, the key entity and entity types of “xrays” and “test” along with the answer types and question type indicates that the question is asking for the reason for the x-ray test and that the answer is a passage from clinical notes.

[0049] The question profile further comprises a number of answers characteristic which indicates whether the answer to the question is to be found in a single clinical note or across multiple clinical notes or passages. In addition, the question profile comprises an answer clinical note type (answer note type) which indicates the type of clinical note in which the answer is to be found for this type of question, e.g., lab reports, structured procedures, progress notes, discharge summaries, etc. These are classifications of note types that may be present in a patient EMR and such notes may be tagged with similar classifications in the patient EMR.

[0050] The question profile may further comprise an indication of the distribution in the patient EMR where the answer is likely to be found, e.g., evenly distributed across all recent lab notes, evenly distributed across progress notes and discharge summaries around the time of the lab report that mentions the xray. Moreover, the question profile may comprise a cardinality that indicates the maximum number of passages or clinical notes that likely provide a portion of the answer to the question.

[0051] The cardinality may be used to determine the number of QA pairs that are added to the labeled training set. The other features of the question profile may be used to perform a weighted calculation of the dynamic confidence score threshold value to be used to select QA pairs, such as by way of using a linear combination of the features as a threshold, for example. Thus, logic may be provided that selects a number of QA pairs corresponding to the cardinality of the question profile associated with the unanswered question in the unlabeled set **U 120** that have confidence scores equal to or above the dynamically generated confidence score threshold value. In some cases, cardinality may take priority over confidence score threshold and thus, if there are not enough QA pairs that have a sufficiently high enough confidence score to equal or exceed the threshold,

then the corresponding top highest scoring confidence score QA pairs may be selected. Alternatively, in other illustrative embodiments, only those that are equal to or above the threshold are available for selection and thus, less than the cardinality number of QA pairs may be selected. As noted above, if a cardinality is not specified, then a default number may be utilized.

[0052] In addition, the features set forth in the question profiles may be used as a basis for performing clustering operations for clustering similar questions, i.e. questions that are directed to finding a similar answer. For example, unanswered questions in the unlabeled set may be associated with a question profile of a similar question in the labeled training set, and that associated question profile may then be used to cluster the unanswered question with other questions in the unlabeled set that are associated with the same question profile or other question profiles that have similar features. This may utilize a similarity scoring algorithm that scores the similarities of features of questions profiles and compares these similarity scores to a threshold, with questions having a high enough similarity score (e.g., equal to or above the threshold) being combined into a similar question group or cluster, e.g., the values of features may be used to generate a feature vector and the feature vectors may be compared, such as by using cosine similarity, to generate a semantic similarity score where if the semantic similarity score is equal to or higher than a threshold value, then the questions are considered to be part of the same group or cluster. Based on the clustering performed, the selected QA pairs for a first unanswered question in the cluster may be used to modify the confidence scores of QA pairs for other similar questions in the cluster that include the answer of the selected QA pairs for the first unanswered question. For example, if a first QA pair has question Q1 and answer A1, and a second question Q2 in the same cluster also has a second QA pair with question Q2 and answer A1, the confidence score of that second QA pair may be increased by a weighting factor, which is implementation dependent, since the first QA pair was selected for the first question Q1 for inclusion in the expansion of the labeled training set **L 110**.

[0053] Thus, with the mechanisms of the illustrative embodiments, given a relatively small labeled training set **110** having a set of ground truth question and answer pairs associated with corresponding question profiles, and a relatively larger unlabeled set **U 120** of unanswered questions, an iterative self-training process is performed that alleviates the need for significant human effort to generate large labeled training sets for training the model to be used by a cognitive system when performing cognitive operations. The self-training process trains the machine learning model based on the labeled training set **110** which is iteratively expanded automatically by correlating unanswered questions with every possible answer in a patient’s EMR, e.g., every clinical note in the patient EMR, to generate a plurality of QA pairs for each unanswered question and then selecting those that have the highest confidence of being a correct answer for the unanswered question. That is, the trained model to that point in time is used to generate confidence scores for the QA pairs. The confidence scores of QA pairs may be modified based on clustering with other questions and QA pairs selected for other similar unanswered questions that have been added to the labeled training set **L 110** in a previous iteration. A question profile of a

question in the labeled training set L 110 is associated with the unanswered question based on a similarity analysis and based on the characteristics of the associated question profile, dynamic selection of QA pairs based on their confidence scores is performed with the selected QA pairs being added to the labeled training set L 110. The process may then be iteratively repeated until either all of the unanswered questions in the unlabeled set 120 are added to the labeled training set L 110, or the self-trained machine learning model converges, i.e., the differences between one iteration of training and another is less than a threshold difference.

[0054] From the above, it can be appreciated that the illustrative embodiments may be utilized in many different types of data processing environments. In order to provide a context for the description of the specific elements and functionality of the illustrative embodiments, FIGS. 3-5 are provided hereafter as example environments in which aspects of the illustrative embodiments may be implemented. It should be appreciated that FIGS. 3-5 are only examples and are not intended to assert or imply any limitation with regard to the environments in which aspects or embodiments of the present invention may be implemented. Many modifications to the depicted environments may be made without departing from the spirit and scope of the present invention.

[0055] FIGS. 3-5 are directed to describing an example cognitive system for healthcare applications (also referred to herein as a “healthcare cognitive system”) which implements a request processing pipeline, such as a Question Answering (QA) pipeline (also referred to as a Question/Answer pipeline or Question and Answer pipeline) for example, request processing methodology, and request processing computer program product with which the mechanisms of the illustrative embodiments are implemented. These requests may be provided as structure or unstructured request messages, natural language questions, or any other suitable format for requesting an operation to be performed by the healthcare cognitive system. As described in more detail hereafter, the particular healthcare application that is implemented in the cognitive system of the present invention is a healthcare application for obtaining answers to natural language questions posed by medical personnel which may be answered based on a corpus of medical knowledge documents and/or patient electronic medical records (EMRs). Thus, in one illustrative embodiment, medical personnel may ask natural language questions about a particular patient, and the healthcare cognitive system will provide answers to those questions based on its analysis of the patient’s EMRs using a trained model which is trained using a self-training process in accordance with the illustrative embodiments.

[0056] It should be appreciated that the healthcare cognitive system, while shown as having a single request processing pipeline in the examples hereafter, may in fact have multiple request processing pipelines. Each request processing pipeline may be separately trained and/or configured to process requests associated with different domains or be configured to perform the same or different analysis on input requests (or questions in implementations using a QA pipeline), depending on the desired implementation. For example, in some cases, a first request processing pipeline may be trained to operate on input requests directed to a first medical malady domain (e.g., various types of blood diseases) while another request processing pipeline may be

trained to answer input requests in another medical malady domain (e.g., various types of cancers). In other cases, for example, the request processing pipelines may be configured to provide different types of cognitive functions or support different types of healthcare applications, such as one request processing pipeline being used for patient diagnosis, another request processing pipeline being configured for medical treatment recommendation, another request processing pipeline being configured for patient monitoring, etc.

[0057] Moreover, each request processing pipeline may have their own associated corpus or corpora that they ingest and operate on, e.g., one corpus for blood disease domain documents and another corpus for cancer diagnostics domain related documents in the above examples. In some cases, the request processing pipelines may each operate on the same domain of input questions but may have different configurations, e.g., different annotators or differently trained annotators, such that different analysis and potential answers are generated. The healthcare cognitive system may provide additional logic for routing input questions to the appropriate request processing pipeline, such as based on a determined domain of the input request, combining and evaluating final results generated by the processing performed by multiple request processing pipelines, and other control and interaction logic that facilitates the utilization of multiple request processing pipelines.

[0058] As noted above, one type of request processing pipeline with which the mechanisms of the illustrative embodiments may be utilized is a Question Answering (QA) pipeline. The description of example embodiments of the present invention hereafter will utilize a QA pipeline as an example of a request processing pipeline that may be augmented to include mechanisms in accordance with one or more illustrative embodiments. It should be appreciated that while the present invention will be described in the context of the cognitive system implementing one or more QA pipelines that operate on an input question, the illustrative embodiments are not limited to such. Rather, the mechanisms of the illustrative embodiments may operate on requests that are not posed as “questions” but are formatted as requests for the cognitive system to perform cognitive operations on a specified set of input data using the associated corpus or corpora and the specific configuration information used to configure the cognitive system. For example, rather than asking a natural language question of “What diagnosis applies to patient P?”, the cognitive system may instead receive a request of “generate diagnosis for patient P,” or the like. It should be appreciated that the mechanisms of the QA system pipeline may operate on requests in a similar manner to that of input natural language questions with minor modifications. In fact, in some cases, a request may be converted to a natural language question for processing by the QA system pipelines if desired for the particular implementation.

[0059] As will be discussed in greater detail hereafter, the illustrative embodiments may be integrated in, augment, and extend the functionality of these QA pipeline, or request processing pipeline, mechanisms of a healthcare cognitive system with regard to the training of the models implemented by the pipeline/cognitive system to analyze input natural language questions and/or evaluate the corpus or corpora of documentation used to answer such natural language questions. In particular, the illustrative embodi-

ments provide mechanisms for automatically expanding a relatively small initial set of ground truth questions and answers used to train a machine learning model, by iteratively performing confidence scoring of a larger set of unanswered questions with regard to a set of training data, e.g., one or more patient EMRs, and utilizing question profiles to determine the selection of question and answer (QA) pairs to be added to the ground truth question and answers, which then results in another iteration of training of the machine learning model and expansion of the ground truth set of questions and answers.

[0060] Since the illustrative embodiments are described in the context of a request processing pipeline, which is considered to be a QA pipeline in the description herein for purposes of illustration, it is important to first have an understanding of how cognitive systems and question and answer creation in a cognitive system implementing a QA pipeline is implemented before describing how the mechanisms of the illustrative embodiments are integrated in and augment such cognitive systems and request processing pipeline, or QA pipeline, mechanisms. It should be appreciated that the mechanisms described in FIGS. 3-5 are only examples and are not intended to state or imply any limitation with regard to the type of cognitive system mechanisms with which the illustrative embodiments are implemented. Many modifications to the example cognitive system shown in FIGS. 3-5 may be implemented in various embodiments of the present invention without departing from the spirit and scope of the present invention.

[0061] As an overview, a cognitive system is a specialized computer system, or set of computer systems, configured with hardware and/or software logic (in combination with hardware logic upon which the software executes) to emulate human cognitive functions. These cognitive systems apply human-like characteristics to conveying and manipulating ideas which, when combined with the inherent strengths of digital computing, can solve problems with high accuracy and resilience on a large scale. A cognitive system performs one or more computer-implemented cognitive operations that approximate a human thought process as well as enable people and machines to interact in a more natural manner so as to extend and magnify human expertise and cognition. A cognitive system comprises artificial intelligence logic, such as natural language processing (NLP) based logic, for example, and machine learning logic, which may be provided as specialized hardware, software executed on hardware, or any combination of specialized hardware and software executed on hardware. The logic of the cognitive system implements the cognitive operation(s), examples of which include, but are not limited to, question answering, identification of related concepts within different portions of content in a corpus, intelligent search algorithms, such as Internet web page searches, for example, medical diagnostic and treatment recommendations, and other types of recommendation generation, e.g., items of interest to a particular user, potential new contact recommendations, or the like.

[0062] IBM Watson™ is an example of one such cognitive system which can process human readable language and identify inferences between text passages with human-like high accuracy at speeds far faster than human beings and on a larger scale. In general, such cognitive systems are able to perform the following functions:

- [0063]** Navigate the complexities of human language and understanding
- [0064]** Ingest and process vast amounts of structured and unstructured data
- [0065]** Generate and evaluate hypothesis
- [0066]** Weigh and evaluate responses that are based only on relevant evidence
- [0067]** Provide situation-specific advice, insights, and guidance
- [0068]** Improve knowledge and learn with each iteration and interaction through machine learning processes
- [0069]** Enable decision making at the point of impact (contextual guidance)
- [0070]** Scale in proportion to the task
- [0071]** Extend and magnify human expertise and cognition
- [0072]** Identify resonating, human-like attributes and traits from natural language
- [0073]** Deduce various language specific or agnostic attributes from natural language
- [0074]** High degree of relevant recollection from data points (images, text, voice) (memorization and recall)
- [0075]** Predict and sense with situational awareness that mimic human cognition based on experiences
- [0076]** Answer questions based on natural language and specific evidence
- [0077]** In one aspect, cognitive systems provide mechanisms for answering questions posed to these cognitive systems using a Question Answering pipeline or system (QA system) and/or process requests which may or may not be posed as natural language questions. The QA pipeline or system is an artificial intelligence application executing on data processing hardware that answers questions pertaining to a given subject-matter domain presented in natural language. The QA pipeline receives inputs from various sources including input over a network, a corpus of electronic documents or other data, data from a content creator, information from one or more content users, and other such inputs from other possible sources of input. Data storage devices store the corpus of data. A content creator creates content in a document for use as part of a corpus of data with the QA pipeline. The document may include any file, text, article, or source of data for use in the QA system. For example, a QA pipeline accesses a body of knowledge about the domain, or subject matter area, e.g., financial domain, medical domain, legal domain, etc., where the body of knowledge (knowledgebase) can be organized in a variety of configurations, e.g., a structured repository of domain-specific information, such as ontologies, or unstructured data related to the domain, or a collection of natural language documents about the domain.
- [0078]** Content users input questions to cognitive system which implements the QA pipeline. The QA pipeline then answers the input questions using the content in the corpus of data by evaluating documents, sections of documents, portions of data in the corpus, or the like. When a process evaluates a given section of a document for semantic content, the process can use a variety of conventions to query such document from the QA pipeline, e.g., sending the query to the QA pipeline as a well-formed question which is then interpreted by the QA pipeline and a response is provided containing one or more answers to the question. Semantic content is content based on the relation between signifiers,

such as words, phrases, signs, and symbols, and what they stand for, their denotation, or connotation. In other words, semantic content is content that interprets an expression, such as by using Natural Language Processing.

[0079] As will be described in greater detail hereafter, the QA pipeline receives an input question, parses the question to extract the major features of the question, uses the extracted features to formulate queries, and then applies those queries to the corpus of data. Based on the application of the queries to the corpus of data, the QA pipeline generates a set of hypotheses, or candidate answers to the input question, by looking across the corpus of data for portions of the corpus of data that have some potential for containing a valuable response to the input question. The QA pipeline then performs deep analysis on the language of the input question and the language used in each of the portions of the corpus of data found during the application of the queries using a variety of reasoning algorithms. There may be hundreds or even thousands of reasoning algorithms applied, each of which performs different analysis, e.g., comparisons, natural language analysis, lexical analysis, or the like, and generates a score. For example, some reasoning algorithms may look at the matching of terms and synonyms within the language of the input question and the found portions of the corpus of data. Other reasoning algorithms may look at temporal or spatial features in the language, while others may evaluate the source of the portion of the corpus of data and evaluate its veracity.

[0080] The scores obtained from the various reasoning algorithms indicate the extent to which the potential response is inferred by the input question based on the specific area of focus of that reasoning algorithm. Each resulting score is then weighted against a statistical model. The statistical model captures how well the reasoning algorithm performed at establishing the inference between two similar passages for a particular domain during the training period of the QA pipeline. The statistical model is used to summarize a level of confidence that the QA pipeline has regarding the evidence that the potential response, i.e. candidate answer, is inferred by the question. This process is repeated for each of the candidate answers until the QA pipeline identifies candidate answers that surface as being significantly stronger than others and thus, generates a final answer, or ranked set of answers, for the input question.

[0081] As mentioned above, QA pipeline mechanisms operate by accessing information from a corpus of data or information (also referred to as a corpus of content), analyzing it, and then generating answer results based on the analysis of this data. Accessing information from a corpus of data typically includes: a database query that answers questions about what is in a collection of structured records, and a search that delivers a collection of document links in response to a query against a collection of unstructured data (text, markup language, etc.). Conventional question answering systems are capable of generating answers based on the corpus of data and the input question, verifying answers to a collection of questions for the corpus of data, correcting errors in digital text using a corpus of data, and selecting answers to questions from a pool of potential answers, i.e. candidate answers.

[0082] Content creators, such as article authors, electronic document creators, web page authors, document database creators, and the like, determine use cases for products, solutions, and services described in such content before

writing their content. Consequently, the content creators know what questions the content is intended to answer in a particular topic addressed by the content. Categorizing the questions, such as in terms of roles, type of information, tasks, or the like, associated with the question, in each document of a corpus of data allows the QA pipeline to more quickly and efficiently identify documents containing content related to a specific query. The content may also answer other questions that the content creator did not contemplate that may be useful to content users. The questions and answers may be verified by the content creator to be contained in the content for a given document. These capabilities contribute to improved accuracy, system performance, machine learning, and confidence of the QA pipeline. Content creators, automated tools, or the like, annotate or otherwise generate metadata for providing information useable by the QA pipeline to identify these question and answer attributes of the content.

[0083] Operating on such content, the QA pipeline generates answers for input questions using a plurality of intensive analysis mechanisms which evaluate the content to identify the most probable answers, i.e. candidate answers, for the input question. The most probable answers are output as a ranked listing of candidate answers ranked according to their relative scores or confidence measures calculated during evaluation of the candidate answers, as a single final answer having a highest ranking score or confidence measure, or which is a best match to the input question, or a combination of ranked listing and final answer.

[0084] FIG. 3 depicts a schematic diagram of one illustrative embodiment of a cognitive system 300 implementing a request processing pipeline 308, which in some embodiments may be a question answering (QA) pipeline, in a computer network 302. For purposes of the present description, it will be assumed that the request processing pipeline 308 is implemented as a QA pipeline that operates on structured and/or unstructured requests in the form of input questions. One example of a question processing operation which may be used in conjunction with the principles described herein is described in U.S. Patent Application Publication No. 2011/0125734, which is herein incorporated by reference in its entirety. The cognitive system 300 is implemented on one or more computing devices 304, 305, 310, and 312 (comprising one or more processors and one or more memories, and potentially any other computing device elements generally known in the art including buses, storage devices, communication interfaces, and the like) connected to the computer network 302. The network 302 includes multiple computing devices 304, 305, 310, and 312 in communication with each other and with other devices or components via one or more wired and/or wireless data communication links, where each communication link comprises one or more of wires, routers, switches, transmitters, receivers, or the like. The cognitive system 300, which may be implemented in a server computing device such as server 305, and network 302 enable question processing and answer generation (QA) functionality for one or more cognitive system users via their respective computing devices 310-312. Other embodiments of the cognitive system 300 may be used with components, systems, sub-systems, and/or devices other than those that are depicted herein.

[0085] The cognitive system 300 is configured to implement a QA pipeline 308 that receive inputs from various sources. For example, the cognitive system 300 receives

input from the network 302, one or more corpora of electronic documents 306 and/or 340, cognitive system users, and/or other data and other possible sources of input. In one embodiment, some or all of the inputs to the cognitive system 300 are routed through the network 302. The various computing devices 304 on the network 302 include access points for content creators and QA system users. Some of the computing devices 304 include devices for a database storing the corpus of data 106 (which is shown as a separate entity in FIG. 3 for illustrative purposes only). Portions of the corpus of data 306 may also be provided on one or more other network attached storage devices, in one or more databases, or other computing devices not explicitly shown in FIG. 3. The network 302 includes local network connections and remote connections in various embodiments, such that the cognitive system 300 may operate in environments of any size, including local and global, e.g., the Internet.

[0086] In one embodiment, the content creator creates content in a document of the corpus of data 306 for use as part of a corpus of data with the cognitive system 300. The document includes any file, text, article, or source of data for use in the cognitive system 300. QA system users access the cognitive system 300 via a network connection or an Internet connection to the network 302, and input questions to the cognitive system 300 that are answered by the content in the corpus of data 306. In one embodiment, the questions are formed using natural language. The cognitive system 300 parses and interprets the question via a QA pipeline 308, and provides a response to the cognitive system user, e.g., cognitive system user 310, containing one or more answers to the question. In some embodiments, the cognitive system 300 provides a response to users in a ranked list of candidate answers while in other illustrative embodiments, the cognitive system 300 provides a single final answer or a combination of a final answer and ranked listing of other candidate answers.

[0087] The cognitive system 300 implements the QA pipeline 308 which comprises a plurality of stages for processing an input question and the corpus of data 306. The QA pipeline 308 generates answers for the input question based on the processing of the input question and the corpus of data 306.

[0088] In some illustrative embodiments, the cognitive system 300 may be the IBM Watson™ cognitive system available from International Business Machines Corporation of Armonk, N.Y., which is augmented with the mechanisms of the illustrative embodiments described hereafter. As outlined previously, a QA pipeline of the IBM Watson™ cognitive system receives an input question which it then parses to extract the major features of the question, which in turn are then used to formulate queries that are applied to the corpus of data. Based on the application of the queries to the corpus of data, a set of hypotheses, or candidate answers to the input question, are generated by looking across the corpus of data for portions of the corpus of data that have some potential for containing a valuable response to the input question. The QA pipeline of the IBM Watson™ cognitive system then performs deep analysis on the language of the input question and the language used in each of the portions of the corpus of data found during the application of the queries using a variety of reasoning algorithms.

[0089] The scores obtained from the various reasoning algorithms are then weighted against a statistical model that summarizes a level of confidence that the QA pipeline of the

IBM Watson™ cognitive system has regarding the evidence that the potential response, i.e. candidate answer, is inferred by the question. This process is repeated for each of the candidate answers to generate ranked listing of candidate answers which may then be presented to the user that submitted the input question, or from which a final answer is selected and presented to the user. More information about the QA pipeline of the IBM Watson™ cognitive system may be obtained, for example, from the IBM Corporation website, IBM Redbooks, and the like. For example, information about the QA pipeline of the IBM Watson™ cognitive system can be found in Yuan et al., "Watson and Healthcare," IBM developerWorks, 2011 and "The Era of Cognitive Systems: An Inside Look at IBM Watson and How it Works" by Rob High, IBM Redbooks, 2012.

[0090] As noted above, while the input to the cognitive system 300 from a client device may be posed in the form of a natural language question, the illustrative embodiments are not limited to such. Rather, the input question may in fact be formatted or structured as any suitable type of request which may be parsed and analyzed using structured and/or unstructured input analysis, including but not limited to the natural language parsing and analysis mechanisms of a cognitive system such as IBM Watson™, to determine the basis upon which to perform cognitive analysis and providing a result of the cognitive analysis. In the case of a healthcare based cognitive system, this analysis may involve processing patient medical records, medical guidance documentation from one or more corpora, and the like, to provide a healthcare oriented cognitive system result.

[0091] In the context of the present invention, cognitive system 300 may provide a cognitive functionality for assisting with healthcare based operations. For example, depending upon the particular implementation, the healthcare based operations may comprise patient diagnostics, medical treatment recommendation systems, medical practice management systems, personal patient care plan generation and monitoring, patient electronic medical record (EMR) evaluation for various purposes, such as for identifying patients that are suitable for a medical trial or a particular type of medical treatment, or the like. Thus, the cognitive system 300 may be a healthcare cognitive system 300 that operates in the medical or healthcare type domains and which may process requests for such healthcare operations via the request processing pipeline 308 input as either structured or unstructured requests, natural language input questions, or the like. In one illustrative embodiment, the cognitive system 300 is a patient EMR question and answer system which provides functionality for answering questions posed by patients and/or medical personnel with regard to a patient's medical condition and the contents of the patient EMR. Such a QA system may further provide treatment recommendation functionality for recommending treatments for the patient based on an analysis of the patient's medical conditions and information available in one or more medical corpora having documents directed to the treatment of various medical conditions, e.g., medical guideline documents, medical reference documents, pharmaceutical documentation, medical procedure documentation, and the like.

[0092] As shown in FIG. 3, the cognitive system 300 is further augmented, in accordance with the mechanisms of the illustrative embodiments, to include logic implemented in specialized hardware, software executed on hardware, or any combination of specialized hardware and software

executed on hardware, for implementing a question profile based self-training engine 320. The question profile based self-training engine 320 provides the logic for performing self-training of one or more machine learning models 350 that implement the logic for analyzing natural language questions, identifying candidate answers to those natural language questions, scoring the candidate answers, and/or selecting a set of one or more final answers for return to an originator of a natural language question as the answer(s) to the natural language question.

[0093] As shown in FIG. 3, the question profile based self-training engine 320 (hereafter referred to as the “self-training engine 320”) comprises a labeled training set (L) 322, self-training model logic 324, unlabeled set (U) 326, question profiles 328, question profile generator 329, and labeled training set expansion logic 330. The labeled training set (L) 322 comprises a plurality of question and answer (QA) pairs that represent the ground truth for training the model 350. The ground truth indicates a set of questions for which the correct answer is known to be true and is associated with the question in the labeled training set (L) 322. In accordance with the mechanisms of the illustrative embodiments, the labeled training set (L) 322 is initially a relatively small set of QA pairs, e.g., 100 QA pairs, which is then expanded iteratively through the automated mechanisms for selecting QA pairs from the unlabeled set (U) 326.

[0094] The labeled training set (L) 322 may initially be generated by one or more subject matter experts (SMEs) that provide the questions and their known correct answers. In addition, the SMEs may provide a question profile for each of the questions in the labeled training set (L) 322, which is stored in the question profiles data structures 328. Alternatively, the questions in the labeled training set (L) may be analyzed, such as by question profile generator 329, using natural language processing techniques which extract the features or characteristics of the question and its corresponding answer and populates the question profile fields automatically based on the results of such natural language processing. The question profile generator 329 may comprise its own natural language processing logic or may utilize natural language processing logic present in the cognitive system 300 to perform such natural language processing and extraction of features/characteristics from the questions and answers.

[0095] The unlabeled set (U) 326 initially comprises a set of unanswered questions that have been generated as candidates for expanding the labeled training set (L) 322. Again, these questions may be natural language questions that are generated by one or more SMEs. As noted above, it is relatively easy for SMEs to generate questions that are representative of the types of questions that a patient and/or medical professional may ask of a QA system. The difficulty is generating correct answers for such questions in an automated manner using the cognitive logic of a QA system. Thus, the creation of the initial unlabeled set (U) 326 does not represent a significant amount of human effort. The unlabeled set (U) 326 may, after initially being provided, be processed by the self-training model logic 324 to associated each unanswered question in the unlabeled set (U) 326 with each portion of a training data set to generate a plurality of question answer (QA) pairs for the unanswered question, the set of QA pairs thereby representing all of the possible sources of answers to the unanswered question. For example, if a training data set has portions 1-1000, then each

unanswered question will have 1000 QA pairs that are to be evaluated. In accordance with some illustrative embodiments used as examples herein, the training data set may be part of the corpus 340 and may comprise one or more patient EMRs for one or more patients, and having portions, e.g., clinical notes, that may serve as answers to natural language questions.

[0096] Given the labeled training set (L) 322, the self-training model logic 324 performs a first iteration to train the machine learning model 350 and thereby generate an initial state of the machine learning model 350. The self-training model logic 324 then applies the trained model 350 to the unlabeled set (U) 326 comprising the QA pairs, where in the context of a medical cognitive system, for example, each QA pair comprises an unanswered question that may be posed by a patient or medical professional paired with a clinical note, or portion of a patient EMR, that may represent the answer to the unanswered question. The application of the currently trained model 350 to the unlabeled set (U) 326 results in the generation of confidence scores associated with each of the QA pairings, where the confidence scores indicate a confidence by the currently trained model 350 that the answer of a QA pair is a correct answer for the unanswered question of the QA pair. This generates, for each unanswered question, a set of answers and confidence scores associated with each of the answers, where the answers are clinical notes in, or portions of, the EMR training data set.

[0097] The labeled training set expansion logic 330 provides the logic for selecting QA pairs from the unlabeled set (U) 326 for inclusion in the labeled training set (L) 322 based on their confidence scores and an association of question profiles with the unanswered questions. The labeled training set expansion logic 330 may also provide logic for performing clustering of unanswered questions based on their associated question profile features/characteristics and/or other key features extracted from the unanswered questions, so as to modify the confidence scores associated with QA pairs based on the selection of QA pairs performed for other unanswered questions in a same cluster.

[0098] For example, the labeled training set expansion logic 330 may associate with each unanswered question (Q) in the QA pairs of the unlabeled set (U) 326 a question profile from the question profile data structures 328, which corresponds to the labeled training set (L) 322. That is, a similarity evaluation is performed by the labeled training set expansion logic 330 for determining a similarity between the unanswered question in the unlabeled set (U) 326 and a question in the labeled training set (L) 322. A similarity score based on a correlation of features may be generated and a highest scoring similar question in the labeled training set (L) 322 is identified for the unanswered question in the unlabeled set (U) 326. The selected similar question in the labeled training set (L) 322 has a corresponding question profile in the question profiles data structures 328 which is then applied to the unanswered question in the unlabeled set (U) 326. The associated question profile is used by the labeled training set expansion logic 330 to dynamically calculate a corresponding confidence score threshold value and/or a number of QA pairs, for an unanswered question, to select from the unlabeled set (U) 326 for inclusion in the labeled training set (L) 322, if any.

[0099] That is, based on the question profile associated with an unanswered question, a confidence score threshold value is calculated that indicates the minimum level of

confidence needed for a QA pair associated with the unanswered question, to be considered for inclusion in the labeled training set (L) 322. It should be appreciated that if sufficient information is not present in the question profile, or if a question profile cannot be identified for the unanswered question, then a default threshold confidence score value may be utilized. Moreover, while this threshold value indicates a minimum level of confidence, in some implementations this minimum level of confidence may be superseded by a required number of QA pairs to be selected such that if there are not a sufficient number of QA pairs that have confidence scores equal to or exceeding the threshold value, some QA pairs that fall below the threshold value may be selected for inclusion in the labeled training set (L) 322.

[0100] In addition, based on information in the question profile associated with the unanswered question, a number of QA pairs to be selected for inclusion in the labeled training set (L) 322 may be identified. This number may be determined based on the cardinality indicated in the question profile, for example. Alternatively, a plurality of features/characteristics specified in the question profile may be evaluated using weighted logic to generate a number of QA pairs to be selected. In general, the number of QA pairs to be selected is combined with the minimum confidence score threshold value such that the number of QA pairs selected comprises a corresponding number of QA pairs selected from those QA pairs that meet or exceed the minimum confidence score threshold. In some illustrative embodiments, the highest scoring QA pairs are selected first prior to selecting lower scoring QA pairs.

[0101] Thus, based on the identified dynamic confidence score threshold, or a default threshold, and the confidence scores associated with the QA pairs in the unlabeled set U 120 for the question, QA pairs that have confidence scores meeting or exceeding the threshold confidence score value are selected by the labeled training set expansion logic 330. The number of such QA pairs that are selected may be based on the dynamic determination corresponding to the question profile as discussed above. These QA pairs are added to the labeled training set (L) 322 to thereby expand the labeled training set (L) 322 to include additional QA pairs as part of the ground truth. The expanded labeled training set (now referred to as L') is then used to update the training of the machine learning model 350 using the expanded QA pairs of the labeled training set (L') 322. This process is repeated until either all of the unanswered questions from the unlabeled set (U) 326 have been added to the labeled training set (L) 322, i.e. at least one QA pair for each unanswered question has been added to the labeled training set (L) 322, or the machine learning model 350 has reached a convergence state, e.g., a change in operational parameters of the machine learning model 350 does not meet or exceed a threshold change level, e.g., no further appreciable change in the operational parameters of the model 350 is performed.

[0102] As noted above, one additional functionality performed by the labeled training set expansion logic 330 is to perform clustering of unanswered questions based on features/characteristics present in their associated question profiles. That is, a clustering operation may be performed on unanswered questions in the unlabeled set U 326 by applying a similarity analysis and scoring logic that scores the degree of similarity between one or more characteristics/features present in the question profiles. The similarity score may be compared to a threshold similarity score value and

if the threshold is met or exceeded, then the two unanswered questions are considered to be sufficiently similar to be clustered together. For instance, a question cluster may contain a set of questions with a high semantic similarity score, i.e. all the questions in the same cluster are essentially asking for the same answer in different ways.

[0103] Based on the clustering, QA pairs selected for inclusion in the labeled training set (L) 322 during a previous iteration of the self-training operation may be used as a basis for modifying confidence scores for similar QA pairs of other unanswered questions in the same cluster. Thus, greater confidence is given to QA pairs having similar answers to that in QA pairs already added to the labeled training set (L) 322 in previous iterations. In this way, different questions that are asking for similar answers are added to the labeled training set (L) 322 to thereby cover various possible ways in which patients and/or medical professionals are likely to ask for particular answers.

[0104] As mentioned above, this process is iteratively repeated with each iteration potentially resulting in a larger expanded labeled training set (L) 322 which is expanded automatically based on the operation of the question profile based self-training engine 320. The resulting trained model(s) 350 may then be output or otherwise provided to the cognitive system 300 which implements these model(s) 350 in the operation of the request processing pipeline 308. The model(s) 350 may be used to analyze the input questions/requests, generate candidate answers/results for the input questions/requests, score the candidate answers/results, and/or select one or more final answers/results to be returned to an originator of the input question/request, e.g., a user of client computing device 310 or 312.

[0105] As is evident from the above, the mechanisms of the illustrative embodiments are rooted in the computer technology arts and are implemented using logic present in such computing or data processing systems. These computing or data processing systems are specifically configured, either through hardware, software, or a combination of hardware and software, to implement the various operations described above. As such, FIG. 4 is provided as an example of one type of data processing system in which aspects of the present invention may be implemented. Many other types of data processing systems may be likewise configured to specifically implement the mechanisms of the illustrative embodiments.

[0106] FIG. 4 is a block diagram of an example data processing system in which aspects of the illustrative embodiments are implemented. Data processing system 400 is an example of a computer, such as server 305 or client 310 in FIG. 3, in which computer usable code or instructions implementing the processes for illustrative embodiments of the present invention are located. In one illustrative embodiment, FIG. 4 represents a server computing device, such as a server 305, which implements a cognitive system 300 and QA system pipeline 308 augmented to include the additional mechanisms of the illustrative embodiments described hereafter.

[0107] In the depicted example, data processing system 400 employs a hub architecture including North Bridge and Memory Controller Hub (NB/MCH) 402 and South Bridge and Input/Output (I/O) Controller Hub (SB/ICH) 204. Processing unit 406, main memory 408, and graphics processor

410 are connected to NB/MCH **402**. Graphics processor **410** is connected to NB/MCH **402** through an accelerated graphics port (AGP).

[0108] In the depicted example, local area network (LAN) adapter **412** connects to SB/ICH **404**. Audio adapter **416**, keyboard and mouse adapter **420**, modem **422**, read only memory (ROM) **424**, hard disk drive (HDD) **426**, CD-ROM drive **430**, universal serial bus (USB) ports and other communication ports **432**, and PCI/PCIe devices **434** connect to SB/ICH **404** through bus **438** and bus **440**. PCI/PCIe devices may include, for example, Ethernet adapters, add-in cards, and PC cards for notebook computers. PCI uses a card bus controller, while PCIe does not. ROM **424** may be, for example, a flash basic input/output system (BIOS).

[0109] HDD **426** and CD-ROM drive **430** connect to SB/ICH **404** through bus **440**. HDD **426** and CD-ROM drive **430** may use, for example, an integrated drive electronics (IDE) or serial advanced technology attachment (SATA) interface. Super I/O (SIO) device **436** is connected to SB/ICH **404**.

[0110] An operating system runs on processing unit **406**. The operating system coordinates and provides control of various components within the data processing system **400** in FIG. **4**. As a client, the operating system is a commercially available operating system such as Microsoft® Windows 10®. An object-oriented programming system, such as the Java™ programming system, may run in conjunction with the operating system and provides calls to the operating system from Java™ programs or applications executing on data processing system **400**.

[0111] As a server, data processing system **400** may be, for example, an IBM® eServer™ System p® computer system, running the Advanced Interactive Executive (AIX®) operating system or the LINUX® operating system. Data processing system **200** may be a symmetric multiprocessor (SMP) system including a plurality of processors in processing unit **406**. Alternatively, a single processor system may be employed.

[0112] Instructions for the operating system, the object-oriented programming system, and applications or programs are located on storage devices, such as HDD **426**, and are loaded into main memory **408** for execution by processing unit **406**. The processes for illustrative embodiments of the present invention are performed by processing unit **406** using computer usable program code, which is located in a memory such as, for example, main memory **408**, ROM **424**, or in one or more peripheral devices **426** and **430**, for example.

[0113] A bus system, such as bus **438** or bus **440** as shown in FIG. **4**, is comprised of one or more buses. Of course, the bus system may be implemented using any type of communication fabric or architecture that provides for a transfer of data between different components or devices attached to the fabric or architecture. A communication unit, such as modem **422** or network adapter **412** of FIG. **4**, includes one or more devices used to transmit and receive data. A memory may be, for example, main memory **408**, ROM **424**, or a cache such as found in NB/MCH **402** in FIG. **4**.

[0114] Those of ordinary skill in the art will appreciate that the hardware depicted in FIGS. **3** and **4** may vary depending on the implementation. Other internal hardware or peripheral devices, such as flash memory, equivalent non-volatile memory, or optical disk drives and the like, may be used in addition to or in place of the hardware depicted

in FIGS. **3** and **4**. Also, the processes of the illustrative embodiments may be applied to a multiprocessor data processing system, other than the SMP system mentioned previously, without departing from the spirit and scope of the present invention.

[0115] Moreover, the data processing system **400** may take the form of any of a number of different data processing systems including client computing devices, server computing devices, a tablet computer, laptop computer, telephone or other communication device, a personal digital assistant (PDA), or the like. In some illustrative examples, data processing system **400** may be a portable computing device that is configured with flash memory to provide non-volatile memory for storing operating system files and/or user-generated data, for example. Essentially, data processing system **400** may be any known or later developed data processing system without architectural limitation.

[0116] FIG. **5** is an example diagram illustrating an interaction of elements of a healthcare cognitive system in accordance with one illustrative embodiment. The example diagram of FIG. **5** depicts an implementation of a healthcare cognitive system **500** that is configured to provide medical treatment recommendations for patients. However, it should be appreciated that this is only an example implementation and other healthcare operations may be implemented in other embodiments of the healthcare cognitive system **500** without departing from the spirit and scope of the present invention.

[0117] Moreover, it should be appreciated that while FIG. **5** depicts the patient **502** and user **506** as human figures, the interactions with and between these entities may be performed using computing devices, medical equipment, and/or the like, such that entities **502** and **506** may in fact be computing devices, e.g., client computing devices. For example, the interactions **504**, **514**, **516**, and **530** between the patient **502** and the user **506** may be performed orally, e.g., a doctor interviewing a patient, and may involve the use of one or more medical instruments, monitoring devices, or the like, to collect information that may be input to the healthcare cognitive system **500** as patient attributes **518**. Interactions between the user **306** and the healthcare cognitive system **500** will be electronic via a user computing device (not shown), such as a client computing device **310** or **312** in FIG. **3**, communicating with the healthcare cognitive system **500** via one or more data communication links and potentially one or more data networks.

[0118] As shown in FIG. **5**, in accordance with one illustrative embodiment, a patient **502** presents symptoms **504** of a medical malady or condition to a user **506**, such as a healthcare practitioner, technician, or the like. The user **506** may interact with the patient **502** via a question **514** and response **516** exchange where the user gathers more information about the patient **502**, the symptoms **504**, and the medical malady or condition of the patient **502**. It should be appreciated that the questions/responses may in fact also represent the user **506** gathering information from the patient **502** using various medical equipment, e.g., blood pressure monitors, thermometers, wearable health and activity monitoring devices associated with the patient such as a FitBit™, a wearable heart monitor, or any other medical equipment that may monitor one or more medical characteristics of the patient **502**. In some cases such medical equipment may be medical equipment typically used in hospitals or medical centers to monitor vital signs and

medical conditions of patients that are present in hospital beds for observation or medical treatment.

[0119] In response, the user **502** submits a request **508** to the healthcare cognitive system **500**, such as via a user interface on a client computing device that is configured to allow users to submit requests to the healthcare cognitive system **500** in a format that the healthcare cognitive system **500** can parse and process. The request **508** may include, or be accompanied with, information identifying patient attributes **518**. These patient attributes **518** may include, for example, an identifier of the patient **502** from which patient EMRs **522** for the patient may be retrieved, demographic information about the patient, the symptoms **504**, and other pertinent information obtained from the responses **516** to the questions **514** or information obtained from medical equipment used to monitor or gather data about the condition of the patient **502**. Any information about the patient **502** that may be relevant to a cognitive evaluation of the patient by the healthcare cognitive system **500** may be included in the request **508** and/or patient attributes **518**.

[0120] The healthcare cognitive system **500** provides a cognitive system that is specifically configured to perform an implementation specific healthcare oriented cognitive operation. In the depicted example, this healthcare oriented cognitive operation is directed to providing a treatment recommendation **528** to the user **506** to assist the user **506** in treating the patient **502** based on their reported symptoms **504** and other information gathered about the patient **502** via the question **514** and response **516** process and/or medical equipment monitoring/data gathering. The healthcare cognitive system **500** operates on the request **508** and patient attributes **518** utilizing information gathered from the medical corpus and other source data **526**, treatment guidance data **524**, and the patient EMRs **522** associated with the patient **502** to generate one or more treatment recommendation **528**. The treatment recommendations **528** may be presented in a ranked ordering with associated supporting evidence, obtained from the patient attributes **518** and data sources **522-526**, indicating the reasoning as to why the treatment recommendation **528** is being provided and why it is ranked in the manner that it is ranked.

[0121] For example, based on the request **508** and the patient attributes **518**, the healthcare cognitive system **500** may operate on the request, such as by using a QA pipeline type processing as described herein, to parse the request **508** and patient attributes **518** to determine what is being requested and the criteria upon which the request is to be generated as identified by the patient attributes **518**, and may perform various operations for generating queries that are sent to the data sources **522-526** to retrieve data, generate candidate treatment recommendations (or answers to the input question), and score these candidate treatment recommendations based on supporting evidence found in the data sources **522-526**. In the depicted example, the patient EMRs **522** is a patient information repository that collects patient data from a variety of sources, e.g., hospitals, laboratories, physicians' offices, health insurance companies, pharmacies, etc. The patient EMRs **522** store various information about individual patients, such as patient **502**, in a manner (structured, unstructured, or a mix of structured and unstructured formats) that the information may be retrieved and processed by the healthcare cognitive system **500**. This patient information may comprise various demographic information about patients, personal contact information about patients,

employment information, health insurance information, laboratory reports, physician reports from office visits, hospital charts, historical information regarding previous diagnoses, symptoms, treatments, prescription information, etc. Based on an identifier of the patient **502**, the patient's corresponding EMRs **522** from this patient repository may be retrieved by the healthcare cognitive system **500** and searched/processed to generate treatment recommendations **528**.

[0122] The treatment guidance data **524** provides a knowledge base of medical knowledge that is used to identify potential treatments for a patient based on the patient's attributes **518** and historical information presented in the patient's EMRs **522**. This treatment guidance data **524** may be obtained from official treatment guidelines and policies issued by medical authorities, e.g., the American Medical Association, may be obtained from widely accepted physician medical and reference texts, e.g., the Physician's Desk Reference, insurance company guidelines, or the like. The treatment guidance data **524** may be provided in any suitable form that may be ingested by the healthcare cognitive system **500** including both structured and unstructured formats.

[0123] In some cases, such treatment guidance data **524** may be provided in the form of rules that indicate the criteria required to be present, and/or required not to be present, for the corresponding treatment to be applicable to a particular patient for treating a particular symptom or medical malady/condition. For example, the treatment guidance data **524** may comprise a treatment recommendation rule that indicates that for a treatment of Decitabine, strict criteria for the use of such a treatment is that the patient **502** is less than or equal to 60 years of age, has acute myeloid leukemia (AML), and no evidence of cardiac disease. Thus, for a patient **502** that is 59 years of age, has AML, and does not have any evidence in their patient attributes **518** or patient EMRs indicating evidence of cardiac disease, the following conditions of the treatment rule exist:

[0124] Age≤60 years=59 (MET);

[0125] Patient has AML=AML (MET); and

[0126] Cardiac Disease=false (MET)

Since all of the criteria of the treatment rule are met by the specific information about this patient **502**, then the treatment of Decitabine is a candidate treatment for consideration for this patient **502**. However, if the patient had been 69 years old, the first criterion would not have been met and the Decitabine treatment would not be a candidate treatment for consideration for this patient **502**. Various potential treatment recommendations may be evaluated by the healthcare cognitive system **500** based on ingested treatment guidance data **524** to identify subsets of candidate treatments for further consideration by the healthcare cognitive system **500** by scoring such candidate treatments based on evidential data obtained from the patient EMRs **522** and medical corpus and other source data **526**.

[0127] For example, data mining processes may be employed to mine the data in sources **522** and **526** to identify evidential data supporting and/or refuting the applicability of the candidate treatments to the particular patient **502** as characterized by the patient's patient attributes **518** and EMRs **522**. For example, for each of the criteria of the treatment rule, the results of the data mining provides a set of evidence that supports giving the treatment in the cases where the criterion is "MET" and in cases where the

criterion is “NOT MET.” The healthcare cognitive system 500 processes the evidence in accordance with various cognitive logic algorithms to generate a confidence score for each candidate treatment recommendation indicating a confidence that the corresponding candidate treatment recommendation is valid for the patient 502. The candidate treatment recommendations may then be ranked according to their confidence scores and presented to the user 506 as a ranked listing of treatment recommendations 528. In some cases, only a highest ranked, or final answer, is returned as the treatment recommendation 528. The treatment recommendation 528 may be presented to the user 506 in a manner that the underlying evidence evaluated by the healthcare cognitive system 500 may be accessible, such as via a drilldown interface, so that the user 506 may identify the reasons why the treatment recommendation 528 is being provided by the healthcare cognitive system 500.

[0128] In accordance with the illustrative embodiments herein, the healthcare cognitive system 500 is augmented to include a question profile self-training engine 540. The question profile self-training engine 540 comprises logic elements 542-550 which operate in a similar manner as previously described above with regard to corresponding elements 322-330 in FIG. 3. The question profile self-training engine 540 provides trained model(s) 560 to the healthcare cognitive system 500 for implementation during runtime processing of requests 508, based on supplied patient attributes 518, to provide treatment recommendations 528. The trained model(s) 560 may be implemented by the healthcare cognitive system 500 to evaluate the various data from corpora 522, 524, and 526 to generate such treatment recommendations 528. Moreover, as shown, the patient EMRs 522 corpus, or a set of one or more patient EMRs for one or more patients in the patient EMR 522 corpus, may be used as a training data set that is input to the question profile based self-training engine 540 to perform training, by the self-training model logic 544, of the model(s) 560 using an iterative self-training process in which the labeled training set (L) 542 is iteratively expanded based on the QA pairs in the unlabeled set (U) 546 and the question profiles 548 and analysis performed by the labeled training set expansion logic 550.

[0129] While FIG. 5 is depicted with an interaction between the patient 502 and a user 506, which may be a healthcare practitioner such as a physician, nurse, physician's assistant, lab technician, or any other healthcare worker, for example, the illustrative embodiments do not require such. Rather, the patient 502 may interact directly with the healthcare cognitive system 500 without having to go through an interaction with the user 506 and the user 506 may interact with the healthcare cognitive system 500 without having to interact with the patient 502. For example, in the first case, the patient 502 may be requesting 508 treatment recommendations 528 from the healthcare cognitive system 500 directly based on the symptoms 504 provided by the patient 502 to the healthcare cognitive system 500. Moreover, the healthcare cognitive system 500 may actually have logic for automatically posing questions 514 to the patient 502 and receiving responses 516 from the patient 502 to assist with data collection for generating treatment recommendations 528. In the latter case, the user 506 may operate based on only information previously gathered and present in the patient EMR 522 by sending a request 508 along with patient attributes 518 and obtaining treatment

recommendations in response from the healthcare cognitive system 500. Thus, the depiction in FIG. 5 is only an example and should not be interpreted as requiring the particular interactions depicted when many modifications may be made without departing from the spirit and scope of the present invention. It should be appreciated, however, that at no time should the treatment itself be administered to the patient 502 without prior approval of the healthcare professional treating the patient, i.e. final determinations as to treatments given to a patient will always fall on the healthcare professional with the mechanisms of the illustrative embodiments serving only as an advisory tool for the healthcare professional (user 506) and/or patient 502.

[0130] As mentioned above, the healthcare cognitive system 500 may include a request processing pipeline, such as request processing pipeline 308 in FIG. 3, which may be implemented, in some illustrative embodiments, as a Question Answering (QA) pipeline. The QA pipeline may receive an input question, such as “what is the appropriate treatment for patient P?”, or a request, such as “diagnose and provide a treatment recommendation for patient P.”

[0131] Thus, the illustrative embodiments provide mechanisms for performing automatic expansion of a labeled training set of QA pairs for improving the training of machine learning models implemented by cognitive systems to perform cognitive operations. The illustrative embodiments significantly reduce the amount of human effort needed to perform such training by provided automated processes that iteratively expand from a relatively small ground truth set of training questions and answers, to obtain a large set of training questions and answers that are more representative of the variety of questions and corresponding answers that a user may ask of a cognitive system in order to obtain a particular answer. This allows the trained model to more accurately handle the various possible questions that may be asked of it when performing question answering operations or otherwise responding to requests for information. Hence, similar result of accurate processing of questions and requests is achieved with significantly less investment of human effort and expenditure of time and resources.

[0132] FIG. 6 is a flowchart outlining an example operation for performing an iterative question profile based self-training operation in accordance with one illustrative embodiment. The operation outlined in FIG. 6 may be implemented, for example, by a question profile self-training engine implemented in one or more computing devices, such as question profile self-training engine 320 in FIG. 3 or 540 in FIG. 5, for example.

[0133] As shown in FIG. 6, the operation starts by receiving an initial labeled training set (L) comprising an initial ground truth set of QA pairs (step 610). A set of unanswered questions is also received (step 620) and an unlabeled set (U) is generated that has QA pairs comprising, for each unanswered question, a set of QA pairs that pair the unanswered question with each possible answer in a training data set (step 630). A machine learning model is trained using the current state of the labeled training set, e.g., either the initial labeled training set L or an expanded labeled training set L' that has been expanded based on a previous iteration of the operation outlined in FIG. 6 (step 640).

[0134] A determination is made as to whether a convergence state has been achieved (step 650). This convergence state may be, for example, that training of the model has not results in an appreciable change in any of the operational

parameters of the model, e.g., equal to or above a threshold amount of change. Convergence may also be determined if all of the unanswered questions have had at least one corresponding QA pair added to the labeled training set L. In some cases, convergence may be a predetermined number of iterations. Other convergence conditions may also be used without departing from the spirit and scope of the present invention.

[0135] If convergence has been achieved, then the trained model is output to the cognitive system for implementation during performance of cognitive operations (step **655**). If convergence has not been achieved, then the trained model is applied to the unlabeled set U to generate confidence scores for the QA pairs (step **660**). Question profiles are associated with the unanswered questions (step **670**) and, for each unanswered question still represented in the unlabeled set U, dynamic selection criteria are generated based on the associated question profiles, e.g., dynamically generated confidence score thresholds and/or number of QA pairs to be selected (step **680**). In some cases, as noted above, this dynamic generation of selection criteria may result in a default being utilized.

[0136] QA pairs are then selected for the unanswered questions based on the selection criteria and the confidence scores associated with the QA pairs (step **690**). The selected QA pairs are added to the labeled training set L to generate an expanded labeled training set L' (which is then considered to be the new labeled training set L) (step **700**) and the operation returns to step **640** where the process is repeated. This process is repeated until the convergence condition is encountered (step **650**) as discussed above.

[0137] As noted above, it should be appreciated that the illustrative embodiments may take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software elements. In one example embodiment, the mechanisms of the illustrative embodiments are implemented in software or program code, which includes but is not limited to firmware, resident software, microcode, etc.

[0138] A data processing system suitable for storing and/or executing program code will include at least one processor coupled directly or indirectly to memory elements through a communication bus, such as a system bus, for example. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution. The memory may be of various types including, but not limited to, ROM, PROM, EPROM, EEPROM, DRAM, SRAM, Flash memory, solid state memory, and the like.

[0139] Input/output or I/O devices (including but not limited to keyboards, displays, pointing devices, etc.) can be coupled to the system either directly or through intervening wired or wireless I/O interfaces and/or controllers, or the like. I/O devices may take many different forms other than conventional keyboards, displays, pointing devices, and the like, such as for example communication devices coupled through wired or wireless connections including, but not limited to, smart phones, tablet computers, touch screen devices, voice recognition devices, and the like. Any known or later developed I/O device is intended to be within the scope of the illustrative embodiments.

[0140] Network adapters may also be coupled to the system to enable the data processing system to become coupled to other data processing systems or remote printers or storage devices through intervening private or public networks. Modems, cable modems and Ethernet cards are just a few of the currently available types of network adapters for wired communications. Wireless communication based network adapters may also be utilized including, but not limited to, 802.11 a/b/g/n wireless communication adapters, Bluetooth wireless adapters, and the like. Any known or later developed network adapters are intended to be within the spirit and scope of the present invention.

[0141] The description of the present invention has been presented for purposes of illustration and description, and is not intended to be exhaustive or limited to the invention in the form 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 embodiment was chosen and described in order to best explain the principles of the invention, the practical application, and to enable others of ordinary skill in the art to understand the invention for various embodiments with various modifications as are suited to the particular use contemplated. 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 in a data processing system comprising a processor and a memory, the memory comprising instructions which cause the processor to be configured to implement a self-training engine of a question and answer system, the method comprising:

pairing, by the self-training engine, an unanswered natural language question with portions of an electronic document to generate an unlabeled data set;

training, by the self-training engine, a model based on a labeled data set comprising a finite number of question and answer pair data structures;

applying, by the self-training engine, the model to the unlabeled data set to identify an answer from the portions of the electronic document to the unanswered natural language question;

modifying, by the self-training engine, the labeled data set to include the identified answer and corresponding unanswered natural language question as an additional question and answer pair data structure; and

training, by the self-training engine, the model based on the modified labeled data set.

2. The method of claim 1, wherein the electronic document is a patient electronic medical record (EMR), and wherein the portions are clinical notes in the patient EMR.

3. The method of claim 1, further comprising generating, for each portion of the electronic document paired with the unanswered natural language question in the unlabeled data set, a confidence value based on application of the model to the unlabeled data set.

4. The method of claim 3, further comprising:

retrieving, for the unanswered natural language question in the unlabeled data set, a question profile specifying

characteristics of a question in the labeled data set that corresponds to the unanswered natural language question; and
 dynamically generating a threshold confidence value for the unanswered natural language question based on the characteristics in the question profile.

5. The method of claim 4, wherein retrieving the question profile comprises:

calculating, for each question in the labeled data set, a similarity score measuring a similarity of first features of the question to second features of the unanswered natural language question;

selecting a question from the labeled data set having a highest similarity score; and

retrieving a profile associated with the selected question as the question profile.

6. The method of claim 4, further comprising:

comparing, for each of the portions of the electronic document paired with the unanswered natural language question, a confidence value of the portion to the dynamically generated threshold confidence value; and
 identifying the answer for the unanswered natural language question from a portion of the electronic document whose corresponding confidence value meets a requirement specified by the threshold confidence value.

7. The method of claim 6, wherein identifying the answer for the unanswered natural language question from a portion of the electronic document whose corresponding confidence value meets a requirement specified by the threshold confidence value comprises selecting a plurality of answers from a plurality of portions of the electronic document, wherein a number of answers in the plurality of answers is determined based on the characteristics in the question profile.

8. The method of claim 4, wherein the characteristics of the unanswered natural language question comprise at least one of key entities associated with the unanswered natural language question, relations between key entities associated with the unanswered natural language question, an answer type, or an indicator of whether the answer is likely to be found in a single portion or across multiple portions of an electronic document.

9. The method of claim 4, wherein the question profile further comprises an indicator of all other questions in the unlabeled data set and labeled data set that are similar to the unanswered natural language question, wherein similar questions are questions whose answers are the same answer.

10. The method of claim 4, further comprising:

clustering questions in the unlabeled data set based on comparisons of features of the questions in the unlabeled data set with characteristics of the question profile to generate a question cluster comprising similar questions, wherein similar questions are questions whose answers are the same answer; and

propagating a label associated with a question answer pair of a first question in the question cluster to all other questions in the question cluster.

11. A computer program product comprising a computer readable storage medium having a computer readable program stored therein, wherein the computer readable program, when executed on a computing device, causes the computing device to:

pair an unanswered natural language question with portions of an electronic document to generate an unlabeled data set;

train a model based on a labeled data set comprising a finite number of question and answer pair data structures;

apply the model to the unlabeled data set to identify an answer from the portions of the electronic document to the unanswered natural language question;

modify the labeled data set to include the identified answer and corresponding unanswered natural language question as an additional question and answer pair data structure; and

train the model based on the modified labeled data set.

12. The computer program product of claim 11, wherein the electronic document is a patient electronic medical record (EMR), and wherein the portions are clinical notes in the patient EMR.

13. The computer program product of claim 11, wherein the computer readable program further causes the computing device to generate, for each portion of the electronic document paired with the unanswered natural language question in the unlabeled data set, a confidence value based on application of the model to the unlabeled data set.

14. The computer program product of claim 13, wherein the computer readable program further causes the computing device to:

retrieve, for the unanswered natural language question in the unlabeled data set, a question profile specifying characteristics of a question in the labeled data set that corresponds to the unanswered natural language question; and

dynamically generate a threshold confidence value for the unanswered natural language question based on the characteristics in the question profile.

15. The computer program product of claim 14, wherein the computer readable program further causes the computing device to retrieve the question profile at least by:

calculating, for each question in the labeled data set, a similarity score measuring a similarity of first features of the question to second features of the unanswered natural language question;

selecting a question from the labeled data set having a highest similarity score; and

retrieving a profile associated with the selected question as the question profile.

16. The computer program product of claim 14, wherein the computer readable program further causes the computing device to:

compare, for each of the portions of the electronic document paired with the unanswered natural language question, a confidence value of the portion to the dynamically generated threshold confidence value; and
 identify the answer for the unanswered natural language question from a portion of the electronic document whose corresponding confidence value meets a requirement specified by the threshold confidence value.

17. The computer program product of claim 16, wherein the computer readable program further causes the computing device to identify the answer for the unanswered natural language question from a portion of the electronic document whose corresponding confidence value meets a requirement

specified by the threshold confidence value at least by selecting a plurality of answers from a plurality of portions of the electronic document, wherein a number of answers in the plurality of answers is determined based on the characteristics in the question profile.

18. The computer program product of claim **14**, wherein the characteristics of the unanswered natural language question comprise at least one of key entities associated with the unanswered natural language question, relations between key entities associated with the unanswered natural language question, an answer type, or an indicator of whether the answer is likely to be found in a single portion or across multiple portions of an electronic document.

19. The computer program product of claim **14**, wherein the question profile further comprises an indicator of all other questions in the unlabeled data set and labeled data set that are similar to the unanswered natural language question, wherein similar questions are questions whose answers are the same answer.

20. An apparatus comprising:
a processor; and
a memory coupled to the processor, wherein the memory comprises instructions which, when executed by the processor, cause the processor to:
pair an unanswered natural language question with portions of an electronic document to generate an unlabeled data set;
train a model based on a labeled data set comprising a finite number of question and answer pair data structures;
apply the model to the unlabeled data set to identify an answer from the portions of the electronic document to the unanswered natural language question;
modify the labeled data set to include the identified answer and corresponding unanswered natural language question as an additional question and answer pair data structure; and
train the model based on the modified labeled data set.

* * * * *