



(51) International Patent Classification:
G06F 17/27 (2006.01)

(21) International Application Number:
PCT/IB2019/055380

(22) International Filing Date:
26 June 2019 (26.06.2019)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
16/022,898 29 June 2018 (29.06.2018) US

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

(71) Applicants (for MG only): **IBM UNITED KINGDOM LIMITED** [GB/GB]; PO Box 41, North Harbour,

Portsmouth Hampshire PO6 3AU (GB). **IBM (CHINA) INVESTMENT COMPANY LIMITED** [CN/CN]; 25/F, Pangu Plaza, No.27, Central North 4th Ring Road, Chaoyang District, Beijing 100101 (CN).

(72) Inventors: **MCATEER, Seamus**; IBM Ireland Product Distribution Limited, Technology Campus, Damastown Industrial Estate, Mulhuddart, 15 (IE). **SALEM, Ahmed**; IBM Ireland Product Distribution Limited, Technology Campus, Damastown Industrial Estate, Mulhuddart, 15 (IE). **MCCLOSKEY, Daniel**; IBM Ireland Product Distribution Limited, Technology Campus, Damastown Industrial Estate, Mulhuddart, 15 (IE). **SOGRIN, Mikhail**; IBM Ireland Product Distribution Limited, Technology Campus, Damastown Industrial Estate, Mulhuddart, 15 (IE).

(74) Agent: **PYECROFT, Justine**; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).

(54) Title: QUERY EXPANSION USING A GRAPH OF QUESTION AND ANSWER VOCABULARY




FIG. 4

(57) Abstract: A method and system are provided for query expansion. The method may include: providing a graph of question and answer word nodes generated from a set of training data for a given domain in the form of training question and answer texts, wherein the answer word nodes are disjoint words that do not occur in both a training question and an associated training answer and wherein edges are provided between a disjoint pair of a training question word and an associated training disjoint answer word, including providing weightings for the nodes and edges based on frequency data; and receiving a user query input, activating input nodes in the graph for words in the user query input, and applying spreading activation through the graph using the weightings to result in a top n most highly activated nodes that are used as candidate words for expansion of the user query input.

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

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

QUERY EXPANSION USING A GRAPH OF QUESTION AND ANSWER VOCABULARY

BACKGROUND

[0001] The present disclosure relates to query expansion in information retrieval, and more specifically, to query expansion using a graph of question and answer vocabulary.

[0002] Information retrieval is the activity of obtaining information resources relevant to an information need from a collection of information resources. Searches can be based on full-text or other content-based indexing, including semantic labeling or other metadata not found directly in the text.

[0003] An information retrieval process begins when a user enters a query into the system. Queries are formal statements of information needs, for example, search strings in search engines. In information retrieval, a query does not uniquely identify a single object in the collection. Instead, several objects may match the query with different degrees of relevancy.

[0004] Query expansion is essential in information retrieval systems. A user should not be expected to know the exact content of documents they hope to retrieve via search. By augmenting a query with additional related terms, the likelihood that a relevant document is retrieved will be increased.

[0005] Query expansion is the process of reformulating a given query to improve retrieval performance in information retrieval operations, particularly in the context of query understanding. In search engines, query expansion involves evaluating a user's input (what words were typed into the search query area, and sometimes other types of data) and expanding the search query to match additional documents.

[0006] Known query expansion involves techniques such as: finding synonyms of words and including these in the search query, finding different morphological forms of the words in the search query, correcting spelling errors, etc. Additional metadata such as related topics of interest or semantic labeling can be used to improve the likelihood that relevant documents are returned, even when a direct lexical or word match is not possible. For example, if words describing cats and dogs are all labeled with the tag "Animal", and a user query contains such vocabulary and is tagged similarly, it may help to filter the set of search results in system output to those that are tagged as containing an "Animal" reference, rather than depending on a direct reference of the specific word in the question.

[0007] Predicting the content of a document that a user's query should point to is far from trivial. Simply adding additional terms to the query may exacerbate the retrieval problem.

[0008] In some situations, a user's query text does not match that of the expected answer. This is particularly exacerbated in domain or task specific searches, such as legal, financial, or biomedical question answering, or in any domain where there is a rich domain specific vocabulary that does not overlap readily with common vocabulary that a typical user may try to search with.

SUMMARY

[0009] According to an aspect of the present disclosure there is provided a computer-implemented method for query expansion using a graph of question and answer vocabulary. The method can include providing a graph of question word nodes and answer word nodes generated from a set of training data for a given domain in the form of training question texts and training answer texts, wherein the answer word nodes are of disjoint words that do not occur in both a training question text and an associated training answer text, and edges are provided between a disjoint pair of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer, including providing weightings for the nodes and edges based on frequency data of the words in the set of training data. The method can further include receiving a user query input, activating input nodes in the graph for at least some words in the user query input, and applying spreading activation through the graph using the weightings to result in a top n most highly activated nodes that are output as candidate words for expansion of the user query input.

[0010] This method provides the advantage that using disjoint answer word nodes in the graph with edges to question word nodes provides disjoint pairs of question-answer words that give a linkage between a set of terms from questions representing a domain task and terms in search content without relying on semantic relationships between the nodes.

[0011] Applying spreading activation may propagate a signal in all directions across question word nodes and answer word nodes that are directly or indirectly connected to the input nodes.

[0012] The method may capture user feedback on results of a user query input into a search engine using the candidate words for expansion and may use the user feedback to update the graph. This provides a graph update mechanism that is easy to implement by augmenting the nodes, edges, and weightings of the graph.

[0013] A user query input may be a question input and activating input nodes may activate question word nodes. In some embodiments, a user query input includes an answer input in addition to the question input and activating input nodes activates answer word nodes in addition to question word nodes. Where a user query input includes an answer input in addition to a question input, this may provide query expansion by refining search results to a style of answer and the candidate words may be used for refinement of the search results to an answer style.

[0014] Activating the graph may be improved using various methods. Activating input nodes in the graph may activate nodes with overlapping surface forms and lemma forms of the words in the user query input. Activating input nodes in the graph may apply edge relationship definitions for edges between question word and answer word pairs with different weightings for each relationship definition of an edge. Activating input nodes may apply sub-topic relationship definitions for edges between question word and answer word pairs with different weightings for each relationship definition of an edge.

[0015] Providing a graph of question word nodes and answer word nodes may include providing a set of training data in a given domain in the form of training question texts and training answer texts, identifying disjoint answer words in a training answer text that do not occur in an associated training question text, generating a graph by adding nodes to the graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint words in the set of training data. Providing the graph can further include adding edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer.

[0016] According to another aspect of the present disclosure there is provided a computer-implemented method for query expansion using a graph of question and answer vocabulary. The method can include providing a set of training data in a given domain in the form of training question texts and training answer texts, identifying disjoint answer words in a training answer text that do not occur in an associated training question text, adding nodes to a graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint answer words in the set of training data. The computer-implemented method can further include adding edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer, and providing weightings for the nodes and edges based on frequency data of the words in the set of training data.

[0017] This aspect generates a graph of disjoint question and answer vocabulary that may be used in the described activation method.

[0018] Providing weightings for the nodes may weight the nodes by the frequency of occurrence of the question word of a question word node in the set of training data and the frequency of occurrence of the answer word of an answer word node in the set of training data. Providing weightings of edges between nodes may weight an edge according to the frequency of the disjoint question word and answer word pair across the set of training data.

[0019] In some embodiments, the graph may be configured for improvement. This may include only creating edges between question word and answer word pairs with the question words and answer words having lexical or

semantic relationships. This may also provide relationship definitions for edges between question word and answer word pairs and may provide different weightings for each relationship definition of an edge.

[0020] Adding nodes to a graph as question word nodes may include conditionally selecting words from the training question texts according to predefined constraints. The predefined constraints may include one or more of a requirement that the word has a pre-specified part-of-speech, a frequency of the term in a corpus, whether or not a word is a stop word for the domain, whether the word appears in a domain vocabulary or ontology, and whether the word is a question word.

[0021] The training question texts may be authored by real intended users and training answer texts may be texts from a corpus of documents from which answers are retrieved.

[0022] According to another aspect of the present disclosure there is provided a system for query expansion using a graph of question and answer vocabulary. The system can include a processor and a memory configured to provide computer program instructions to the processor to execute the function of the components, the components including a graph generating component for providing a graph of question word nodes and answer word nodes generated from a set of training data for a given domain in the form of training question texts and training answer texts, wherein the answer word nodes are of disjoint words that do not occur in both a training question text and an associated training answer text, and edges are provided between a disjoint pair of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer, including providing weightings for the nodes and edges based on frequency data of the words in the set of training data. The components can further include a graph applying component including a query receiving component for receiving a user query input, a graph input component for activating input nodes in the graph for at least some words in the user query input, a graph activation component for applying spreading activation through the graph using the weightings, and a graph output component for returning a result of a top n most highly activated nodes that are output as candidate words for expansion of the user query input. The system may also include a user feedback component for capturing user feedback on results of a user query input into a search engine using the candidate words for expansion and using the user feedback to update the graph.

[0023] The user query input may be a question input and/or an answer input, and the graph input component may activate question word nodes and/or answer word nodes. Where a user query input includes an answer input in addition to the question input, this may refine search results to a style of answer, and wherein the graph output component may return a result of a top n most highly activated nodes that are used as candidate words for refinement of the search results to an answer style.

[0024] According to another aspect of the present disclosure there is provided a system for query expansion using a graph of question and answer vocabulary. The system can include a processor and a memory configured to provide computer program instructions to the processor to execute the function of the components. The components can include a graph generating component including a training data component for providing a set of training data in a given domain in the form of training question texts and training answer texts, a disjoint word component for identifying disjoint answer words in a training answer text that do not occur in both an associated training question text, a node adding component for adding nodes to a graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint answer words in the set of training data. The components can further include an edge adding component for adding edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer, and a weighting component for providing weightings for the nodes and edges based on frequency data of the words in the set of training data. The components can further include a weighting component which may include a node weighting component for providing weightings for the nodes configured to weight the nodes by the frequency of occurrence of the question word of a question word node in the set of training data and the frequency of occurrence of the answer word of an answer word node in the set of training data, and an edge weighting component for providing weightings of edges between nodes configured to weight an edge according to the frequency of the disjoint question word and answer word pair across the set of training data.

[0025] The system may include a semantic restricting component for providing the graph that only creates edges between question word and answer word pairs with the question words and answer words having lexical or semantic relationships.

[0026] The system may include an edge relationship component for providing relationship definitions for edges between question word and answer word pairs and providing different weightings for each relationship definition of an edge.

[0027] The system may include a feedback receiving component for receiving user feedback on results of a user query input into a search engine using the candidate words for expansion and a graph updating component for applying the user feedback to the graph.

[0028] According to another aspect of the present disclosure there is provided a computer program product for query expansion using a graph of question and answer vocabulary, the computer program product including a computer readable storage medium having program instructions embodied therewith. The program instructions executable by a processor can cause the processor to provide a graph of question word nodes and answer word nodes generated from a set of training data for a given domain in the form of training question texts and training answer texts, wherein the answer word nodes are of disjoint words that do not occur in both a training question text

and an associated training answer text and edges are provided between a disjoint pair of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer, including providing weightings for the nodes, and edges based on frequency data of the words in the set of training data. The program instructions can be further configured to cause the processor to receive a user query input, activating input nodes in the graph for at least some words in the user query input, and applying spreading activation through the graph using the weightings to result in a top n most highly activated nodes that are used as candidate words for expansion of the user query input.

[0029] According to another aspect of the present disclosure there is provided a computer program product for query expansion using a graph of question and answer vocabulary, the computer program product including a computer readable storage medium having program instructions embodied therewith. The program instructions executable by a processor to can cause the processor to provide a set of training data in a given domain in the form of training question texts and training answer texts, identify disjoint answer words in a training answer text that do not occur in an associated training question text, add nodes to a graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint answer words in the set of training data. The program instructions can be further configured to cause the processor to add edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer, and provide weightings for the nodes and edges based on frequency data of the words in the set of training data.

BRIEF DESCRIPTION OF THE DRAWINGS

[0030] The drawings included in the present disclosure are incorporated into, and form part of, the specification. They illustrate embodiments of the present disclosure and, along with the description, serve to explain the principles of the disclosure. The drawings are only illustrative of typical embodiments and do not limit the disclosure.

[0031] Some embodiments of the present invention will now be described, by way of example only, with reference to the following drawings in which:

Figure 1 is a block diagram of an example embodiment of a system in which aspects of the present invention are implemented;

Figures 2A-2C are schematic diagrams of an example embodiment illustrating generating a graph in accordance with the present disclosure;

Figure 3 is a flow diagram of an example embodiment of an aspect of generating a graph for use in query expansion of a method in accordance with the present disclosure;

Figure 4 is a flow diagram of an example embodiment of another aspect of applying a query expansion graph in query expansion of a method in accordance with the present disclosure;

Figure 5 is a block diagram of an example embodiment of a query expansion graph generating system in accordance with the present disclosure;

Figure 6 is a block diagram of an example embodiment of a query expansion graph applying system in accordance with the present disclosure;

Figure 7 is a schematic diagram illustrating a graph for query expansion in accordance with an embodiment of the present disclosure;

Figure 8 is a block diagram of an embodiment of a computer system or cloud server in which the present disclosure may be implemented;

Figure 9 is a schematic diagram of a cloud computing environment in which the present disclosure may be implemented; and

Figure 10 is a diagram of abstraction model layers of a cloud computing environment in which the present disclosure may be implemented.

[0032] It will be appreciated that for simplicity and clarity of illustration, elements shown in the figures have not necessarily been drawn to scale. For example, the dimensions of some of the elements may be exaggerated relative to other elements for clarity.

DETAILED DESCRIPTION

[0033] The described method and system provide query expansion by building a graph of disjoint question-answer word pairs from a ground truth in the form of a set of training data and activating the graph with unseen query vocabulary at runtime. Using the described method and system directly wires into answer content that has no immediately apparent lexical or semantic link with the question text.

[0034] Referring to Figure 1, a block diagram shows an example embodiment of an information retrieval system 100 in which the described method and system may be implemented including a search engine 160 for searching a search index 170 of documents and a query expansion system 110 for expanding a user query for use in the search engine 160 for retrieving documents from the search index 170.

[0035] The described query expansion system 110 includes a graph generating component 120 for generating a question and answer graph 150 from a set of training data 140 for a domain. The training data 140 is a ground truth of information provided by direct input, in this case in the form of training questions with training answers in a specific domain.

[0036] The question and answer graph 150 is a data structure providing a graph of disjoint question and answer vocabulary that has question word nodes and answer word nodes that have connecting edges and node and edge weightings generated from the training data 140 as described further below.

[0037] The query expansion system 110 includes a graph applying component 130 for applying the question and answer graph 150 during a user query in a search engine 160 in order to expand the query terms.

[0038] The graph applying component 130 includes a query receiving component 131 for receiving a user's unseen query input, a graph input component 132 for inputting at least some of the words from the query input into the question and answer graph 150. The graph applying component 130 includes a graph activation component 133 for activating the graph with the activated nodes of the query input and a graph output component 134 for outputting a top n words from the graph activation.

[0039] The graph applying component 130 includes an expanded query output component 135 for outputting the expanded query to the search engine 160 for improving the search of the search index 170.

[0040] The illustrated information retrieval system 100 can improve the quality of relevance for documents retrieved through a search query. By augmenting the query with additional related terms, it can increase the likelihood that relevant documents are retrieved.

[0041] Referring to Figure 2A, a schematic diagram 200 illustrates an example embodiment of the aspect of generating a question and answer graph 150 from a set of training data 140 for a domain.

[0042] The set of training data 140 is a ground truth in the form of a set of question texts 241 with associated answer texts 242. The question texts 241 are made up of question words 211 and the answer texts 242 are made up of answer words 221. Such training data is known from training a question-answer or information retrieval system to improve the rank of correct answers. An example set of training data may include data sets from approximately 140 to 50,000 question-answer pairs, but there is no minimum required or maximum permitted number of examples, in order for this method to be used. As the number of examples provided increases, however, the likelihood of unseen question vocabulary appearing in the graph increases, which will in turn increase the relevancy of the returned top n answer word expansions for the user's input.

[0043] The question texts 241 are authored by real intended users of the system and describe the common types of phrases and vocabulary representative of typical users. The associated answer texts 242 are taken from the corpus of documents from which answers or search results will be retrieved.

[0044] Training a question-answer or information retrieval system relies on subject matter experts to provide these question texts 241 and correct answer texts 242. This is a primary source ground truth data that is available in training enterprise systems. Depending on the technique that uses such data, more or less of this type of data may be used in order to produce improvements in search results. This technique does not have a minimum required or maximum permitted number of examples.

[0045] The question vocabulary 210 in the question texts 241 and the answer vocabulary 220 in the answer texts 242 are evaluated to restrict to disjoint question-answer word pairs. Disjoint question-answer word pairs are pairs of words that do not occur in both a training question text 241 and an associated training answer text 242. In some embodiments, the answer words 221 may be restricted to words that are in an answer text 242 but are not included in the question text 241.

[0046] Referring to Figure 2B, these disjoint question-answer word pairs may be considered in the context of a Venn diagram 245 shown for illustrative purposes. Question vocabulary 210 and answer vocabulary 220 are shown as two separate sets of words that sometimes overlap, shown in the Venn diagram 245 as the overlap 230. The overlap 230 or intersection of these two sets is the “conjoint” vocabulary, and this conjoint vocabulary set is the normal basis for search engines, where documents can be retrieved that directly contain these question words.

[0047] Known query expansion attempts to expand this intersection space, typically use synonyms or hyponyms (e.g., “Labrador” = “dog”) or domain-specific knowledge (e.g., “paracetamol = Tylenol” (Tylenol is a trademark of Johnson & Johnson)), but this process of expansion of the intersection typically relies on traditional semantic resources that describe human understandable or obvious semantic relationships, even including relatively distant relationships like “headache” as an expansion for “paracetamol”, since paracetamol “treats” headaches.

[0048] Disjoint vocabulary, on the other hand, is the sets of words that are in the answer vocabulary 220 but are not in the associated question vocabulary 210. Answer words that are in the overlap 230 are not used for question-answer word pairs. Consequently, disjoint pairs 250 are the exhaustive combinations of the question words 211 and with disjoint answer words 221, between question text 241 and associated answer text 242.

[0049] The nature of these disjoint pairs 250 may not describe obvious semantic relationships, but nevertheless captures some aspect of the user’s intent as expressed in their question, and which is answered by the vocabulary of the relevant answer. These may be considered to be non-obvious semantic relationships and prior art has ignored them. This phenomena can be demonstrated in examples like “delay” and “interval”, which while both are time related, have no easily recognized or commonly extracted semantic relationship, and as such are extremely difficult to leverage as expansions using traditional techniques.

[0050] Referring to Figure 2C, a question and answer graph 150 is built from the words in the set of training data 140 (shown in Figure 2A). Question word nodes 251 may be provided for all the question words 211 in the question texts 241 of the training data 140, or a conditional selection of the question words 211 from the question texts 241 such as question words 211 with a part-of-speech, the frequency of a word in the domain corpus, stop words for the domain, etc.

[0051] Answer word nodes 252 are provided for the disjoint answer words 221 in the answer texts 242 of the training data 140. The disjoint answer words 221 are words that are not also in the vocabulary of the question texts 241.

[0052] Edges 253 are provided in the graph 150 between disjoint word pairs 250 in the form of a question word 211 and a disjoint answer word 221. Although there are no direct edges 253 between two answer word nodes 252, answer word nodes 252 may be only 1 node degree away from each other by virtue of the intervening question word node 251 that is linked to them both.

[0053] Node weightings 261, 262 are provided as the frequency of the word as a question word 211 or an answer word 221 in the set of training data 140.

[0054] Edge weighting 263 is provided as the frequency of the particular disjoint pair 250 across all question-answer data in the set of training data 140. This may be the frequency that a disjoint pair 250 of words is disjoint for each question-answer combination in the ground truth.

[0055] Each node 251, 252 or edge 253 stores an associated value for its weighting. These seed values are the basis for a mathematical formula that is applied as the spreading signal traverses these same nodes and edges during activation.

[0056] The graph 150 is activated for input words of an unseen user question with a signal that will move in all directions across all nodes 251, 252, in order to accumulate on the most relevant nodes in the graph 150 for the selected inputs. The signal is able to move freely across all nodes 251, 252 in the graph 150 that are directly and indirectly connected to the input words.

[0057] In one embodiment, part-of-speech (POS) tagging may be applied to question text 241 and/or answer text 242 in processing the training data 140. Edges may only be created between question and answer words that have the same POS tag (e.g. noun, verb, adjective, etc.). In this way, the graph activation may be forced to traverse same POS edges.

[0058] In some embodiments, when edges are added between question and answer word nodes, there may be additional relationships recognized in addition to the default relationship of "isDisjoint". Other types of edge that may be drawn between nodes may be, for example: "sameLanguage" or "differentLanguage", depending on question and associated answer sharing the same language; "carInsuranceRelated", "vanInsuranceRelated", or "houseInsuranceRelated", depending on the sub-topic of the question and associated answers, as determined by a separate topic classifier, model, or semantic labeling mechanism. When the same topic is recognized in a new user input, these edges may be traversed with different weightings, according to their significance to the input.

[0059] Referring to Figure 3, a flow diagram 300 shows an example embodiment of an aspect of the described method of generating a graph of question and answer vocabulary for use in query expansion.

[0060] The method 300 provides 301 a set of training data in a given domain in the form of training question texts and training answer texts. Training data is provided in the form of a corpus of ground truths for a domain in the form of questions with known correct answers. The questions are authored by real intended users of the system and describe the common types of phrases and vocabulary representative of typical users. Associated answer texts are taken from the corpus of documents from which answers or search results will be retrieved for the targeted domain. Therefore, the answers are highly relevant in context but do not necessarily have any lexical or semantic link with the question text. A relatively small amount of training data is required for a ground truth for a domain. This is also easy to augment with additional question-answer pairs. The training data is targeted to the user's domain and the corpus being searched.

[0061] Disjoint answer words are identified 302 in the answer texts, where disjoint words are words that do not occur in an associated training question text as well as the training answer text. Disjoint pairs of words are provided between a word in a training question text and a word in an associated training answer text. There may be more than one associated training answer text for each question text.

[0062] A graph of disjoint question and answer vocabulary is built, which has nodes for vocabulary words of node types QuestionWord and AnswerWord. The method adds 303 nodes to the graph with question words as QuestionWord nodes and adds disjointed answer words as AnswerWord nodes for the identified disjoint words.

[0063] Words from the question vocabulary of the set of training data may become question word nodes regardless of whether they appear in answers or not. This is to fully capture the vocabulary of the task. In some embodiments, all question words may be included as question word nodes in the graph to ensure that hidden context is included from question text.

[0064] In other embodiments, there may be conditional selection of question words from question texts for inclusion as question word nodes. Such conditional selection may include various features. The features may include, as examples, one or more of: the requirement that the word has some pre-specified set of part-of-speech, the frequency of the term in a corpus, whether or not it is a stop word for the domain, whether the term appears in some domain vocabulary/ontology whether the word is a question word, etc. In another embodiment, the question word nodes may also be limited to disjoint words that are not included in an associated answer text.

[0065] In mathematics, graph theory uses mathematical structures to model pairwise relations between objects. A graph in this context is made up of question word nodes and answer word nodes, which are connected by edges. A graph may be undirected, meaning that there is no distinction between the two vertices associated with each edge, or its edges may be directed from one vertex to another. In this embodiment, the graph is undirected.

[0066] The method 300 may add 304 edges to the graph between a disjoint pair of a question word in a training question and an answer word in an associated training answer. The use of disjoint pairs of question-answer words gives an unintuitive but yet meaningful linkage between a set of terms from inquiries representing a domain task and terms in other content (such as the search corpus) which are not retrievable by any heuristic approach such as employed in known methods.

[0067] The method 300 may provide 305 weightings for the nodes and edges based on frequency data of the words in the set of training data.

[0068] Weightings for the nodes may be provided by weighting the nodes by the frequency of occurrence of the question word of a question word node in the set of training data and the frequency of occurrence of the answer word of an answer word node in the set of training data.

[0069] Weightings of edges between nodes may be provided by weighting an edge according to the frequency of the disjoint question word and answer word pairs across the set of training data.

[0070] There are multiple ways to configure the graph for optimal results of spreading signal through the graph. For example, there may be conditional inclusion at the node level and/or this may be managed through the connectivity on the edges. In some embodiments, including everything by default with the spreading activation as described can give good results on real world tasks.

[0071] In some embodiments, the graph may be generated with only edges provided between question word and answer word pairs having lexical or semantic relationships. For example, the edges may only be created for

question words and answer words having the same part-of-speech. This will restrict activation of the graph to the semantic relationship.

[0072] In some embodiments, the method may provide 306 relationship definitions for edges between question word and answer word pairs in the graph and may be configured to provide different weightings for each relationship definition of an edge.

[0073] The graph may be updated by augmenting the set of training data with additional or replacement questions. The graph may also be updated by receiving user feedback of the results of queries made using expanded queries using the graph as described further below.

[0074] Referring to Figure 4, a flow diagram 400 shows an example embodiment of an aspect of the described method of expanding a query using a graph of question and answer vocabulary.

[0075] The method 400 may provide 401 a graph of question word nodes and answer word nodes generated for a set of training data for a given domain as generated in the method of Figure 3. The graph includes question word nodes and answer word nodes. The answer word nodes are of disjoint words that do not occur in both a training question text and an associated training answer text. Edges are provided between disjoint pairs of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer. The graph includes weightings for the nodes and edges based on frequency data of the words in the set of training data.

[0076] The method 400 may receive 402 a user query input in the form of a set of words. A user query input may commonly be a question input; however, the query input may also include an answer input. The input nodes may be nodes with overlapping surface forms or lemma forms of the words in the user query input.

[0077] The method 400 may activate 403 input nodes in the graph for words in the user query input. The graph is activated using user input only. This can be in the form of a question, or a question and an associated answer for further context. Depending on the nature of the input, the appropriate nodes will be activated. If the query input is a question input, then activating input nodes activates question word nodes in the graph. If the query input also includes an answer input, activating input nodes may activate answer word nodes in the graph in addition to question word nodes.

[0078] The method 400 does not assume that the relationships between nodes would be humanly obvious, but that collectively, following the activation of the graph, there is a representation of the user intent, even if it is not immediately apparent. For this reason, all vocabulary of a user query input may be included in the graph by default

without trying to assume that particular words are more or less helpful, initially. The activation of the graph and the distribution of weight across these connections is relied upon to determine if a given word is relevant or not.

[0079] The non-obvious relationships between vocabulary in questions and answers is an interesting factor in the described method 400. A word like "and" or a possessive marker "s" may be critical in a particular context, and it is not assumed that they should be removed simply because they are not obviously related. For example, the general nature of answers about capital cities may be to refer to multiple characteristics collectively, separated by the word "and", and this is something that is possible to capture in a graph of the present disclosure. The benefits are clear in the use of this technique in domains such as legal, insurance, medical, financial, etc.

[0080] The method 400 may apply 404 spreading activation through the graph using the weightings to result in a top n most highly activated nodes. Applying spreading activation propagates a signal in all directions across question word nodes and answer word nodes that are directly or indirectly connected to the input nodes. The spreading activation weight improves information retrieval by allowing context sensitive ranking.

[0081] Lexical and semantic constraints may be applied when applying the spreading activation to expand using same lexical or semantic constraints only, such as the same part-of-speech only.

[0082] In some embodiments, when edges are added between question and answer words, there may be additional relationships recognized in addition to the default relationship of being disjoint pairs. Other types of edge that may be drawn between nodes may be: "sameLanguage" or "differentLanguage", depending on question and associated answer sharing the same language; "carInsuranceRelated", "vanInsuranceRelated", or "houseInsuranceRelated", depending on the sub-topic of the question and associated answers. When the same topic is recognized in a new user input, these edges may be traversed with different weight according to their significance in the input.

[0083] The method 400 uses 405 the n most highly activated nodes as candidate words for expansion of the user query input. This provides context relevant top n words for expansion. A search query may be augmented with most relevant expansion candidates.

[0084] The method 400 may select the top n question or answer word nodes. However, in some embodiments, only answer word nodes are selected as these have the advantage of leveraging answer words that are guaranteed to exist in the background retrieval corpus, whereas this is not always true for question words.

[0085] In the case of a user query input including an answer input in addition to a question input, this query expansion may refine search results to a style of answer, for which a sample answer may be provided. In this case,

both QuestionWord and AnswerWord nodes can be activated at first. After activation, the signal may traverse the graph from those points and accumulate at other nodes in the graph, which is then used as a relevance measure.

[0086] The method 400 may include capturing 406 user feedback on results of a user query input into a search engine using the candidate words for expansion and using the user feedback to supplement the set of training data.

[0087] A search system in which an expanded query generated by the described method 400 is entered may capture user feedback, such as a thumbs rating or a star rating on an answer to a question. This feedback data may be used to supplement the initial ground truth question-answer data, and thereby improve or adapt the system over time.

[0088] Since the initial ground truth is authored by experts, it is normally considered a gold standard and more reliable than raw user feedback. Therefore, the use of the user feedback loop to improve a search system is can be controlled, for example, by monitoring particular trusted authors only, or applying a voting system and a threshold of trustworthiness before inclusion. This use of feedback to improve the system can be a much cheaper method of integrating user data than other techniques for query expansion.

[0089] The frequencies of the nodes and edges in the graph may be augmented with the additional counts from the provided feedback data, so new nodes may be added for any previously unseen question or answer words and/or new edges may be drawn between question and answer pairs that were not seen in previous ground truth data. This process of iterative updating of the graph can be very cheap and computationally efficient to perform as the associated actions in updating the graph do not require a deep retraining or recalculation of feature vectors or other data as is the case in other techniques, such as deep learning methods.

[0090] Once a user gives a positive rating, such as a thumbs up or a high star rating, to a particular answer or answers in a search result, this answer and its associated question can be sent to the underlying query expansion graph and the data used as a supplementary question-answer data to augment the data already there with additional nodes, edges, and/or frequency counts, as described above.

[0091] The graph provides a network of possible connections between vocabulary that, when activated with content of an unseen question, can predict the best set of expansions that existed in correct answers in the set of vocabulary where there was no overlap between the question and the answer. This provides user driven query expansion through activation of disjoint question-answer pairs.

[0092] The output is the most relevant to all the sub-elements of the user query. Using an expanded query as described, the search for relevant documents can be improved. This is due to the focus on contextual gaps in user knowledge or vocabulary. The described method 400 produces statistically relevant expansions as distinct from semantic expansions. The expansions reflect the semantics of user intent rather than content semantics. The method 400 leverages ambient context of the mappings extracted from the user-authored ground truth. Additional user content can result in more accurate representations of context.

[0093] Connections like [mortgage -> build] or [delay -> interval] are now possible even though they do not have the same part of speech, are not synonyms, are not easily discoverable by existing methods, and do not overlap between contexts provided by user data (e.g., queries and answers).

[0094] Although the method 400 may depend on the hard-coded nature of question-answer pairs from a ground truth source, this is much more flexible than a method of hardcoding explicit expansion pairs.

[0095] An improved solution to this problem should increase the relevance of expansions by generating relevant expansions only and by not producing expansions that do not exist in the target corpus. An improved solution should also reduce over-inflation of recall by having an effective ranking of top n expansions.

[0096] The described method 400 allows for easy adaptation to new domains or user contexts by updating the set of training data. The training data may also be used to target expansions that solve inherent gaps in user knowledge by linking user question vocabulary to correct answers.

[0097] Referring to Figure 5, a block diagram shows an example embodiment of the described query expansion system 110 with components providing the functionality of the graph generating component 120.

[0098] In the illustrated embodiment, the query expansion system 110 includes at least one processor 501 (or a hardware module or a circuit) for executing the functions of the described components, which may be software units executing on processor 501. In some embodiments, multiple processors 501 running parallel processing threads may be provided, enabling parallel processing of some or all of the functions of the components. Memory 502 may be configured to provide computer instructions 503 to the at least one processor 501 to carry out the functionality of the components, such as the graph generating component 120 depicted in Figure 5 and the graph applying component 130 depicted in Figure 6.

[0099] In the illustrated embodiment, the graph generating component 120 includes a training data component 522, for providing a set of training data in a given domain in the form of training question texts and

training answer texts, and a disjoint word component 523, for identifying disjoint answer words in a training answer text that do not occur in an associated training question text.

[0100] In the illustrated embodiment, the graph generating component 120 includes a node adding component 524 for adding nodes to a graph as question word nodes and answer word nodes. The answer word nodes may be for the identified disjoint answer words in the set of training data. The graph generating component 120 may include an edge adding component 525 for adding edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer.

[0101] In the illustrated embodiment, the graph generating component 120 includes a frequency determining component 526, for determining frequency data of words and disjoint word pairs in the set of training data, and a weighting component 527, for providing weightings for the nodes and edges based on frequency data of the words in the set of training data. The weighting component 527 includes a node weighting component 528 for providing weightings for the nodes, the node weighting component 528 being configured to weight the nodes by the frequency of occurrence of the question word of a question word node in the set of training data and the frequency of occurrence of the answer word of an answer word node in the set of training data. The weighting component 527 also includes an edge weighting component 529 for providing weightings of edges between nodes configured to weight an edge according to the frequency of the disjoint question word and answer word pair across the set of training data.

[0102] In the illustrated embodiment, the graph generating component 120 includes a semantic restricting component 530 for providing the graph with edges restricted between question word and answer word pairs to question words and answer words having lexical or semantic relationships. The graph generating component 120 also includes an edge relationship component 531 for providing relationship definitions for edges between question word and answer word pairs and providing different weightings for each relationship definition of an edge. The graph generating component 120 also includes a feedback receiving component 532 for receiving user feedback on results of a user query input into a search engine using the candidate words for expansion, and a graph updating component 533 for applying the user feedback to the graph.

[0103] Referring to Figure 6, a block diagram shows an example embodiment of the described query expansion system 110 with components providing the functionality of the graph applying component 130.

[0104] In the illustrated embodiment, the graph applying component 130 includes a query receiving component 131 for receiving a user query input and a graph input component 132 for activating input nodes in the graph for words in the user query input by a node selection component 633. The user query input may be a

question input and/or an answer input, and the graph input component 132 may activate question word nodes and/or answer word nodes.

[0105] In the illustrated embodiment, the graph applying component 130 includes a graph activation component 133, for applying spreading activation through the graph from the input nodes using the weightings of the nodes and edges, and a graph output component 134, for returning a result of a top n most highly activated nodes by an n node output component 637. The words of the output nodes may be used in an expanded query output component 135 as candidate words for expansion of the user query input.

[0106] If a user query input includes an answer input in addition to a question input, this may refine search results to a style of answer. In this case, the graph output component 134 may return a result of a top n most highly activated nodes that are used as candidate words for refinement of the search results to an answer style.

[0107] In the illustrated embodiment, the graph activation component 133 includes a semantic restriction applying component 641 wherein the graph has edges restricted between question word and answer word pairs to question words and answer words having lexical or semantic relationships. The graph activation component 133 also includes an edge relationship applying component 642 for applying provided relationship definitions for edges between question word and answer word pairs and providing different weightings for each relationship definition of an edge.

[0108] In the illustrated embodiment, the query expansion system 110 includes a user feedback component 650 for capturing user feedback on results of a user query input into a search engine using the candidate words for expansion, and the user feedback may be used to supplement the set of training data resulting in an updated graph.

[0109] The described query expansion system 110 including a graph generating component 120 (shown in Figure 5) and a graph applying component 130 may be introduced into existing search systems without re-architecture. This may be provided as part of front-end query building and does not affect a background index. The described query expansion system 100 (shown in Figure 1) may act as a domain-customization interface on any existing search system including open-source systems.

[0110] Example

[0111] Training Data:

[0112] Question-answer data is a map that leads from question text to potentially multiple correct answers. The more such examples that are provided, the better the depth of relationships between vocabulary in the graph, and the better the graph can potentially recommend expansions for a new unseen context. The examples used here are contrived for simplicity in order to demonstrate the mechanism independently of the inherent value of graph depth:

Question: "Is Dublin the capital city of Ireland?"

Answers: ["Ireland's capital and largest city by population is Dublin", "Dublin is the biggest city and capital of Ireland"]

Question: "What type of animal is a lion?"

Answers: ["The lion is the largest cat in Africa.", "The lion is an animal in the cat family."]]

[0113] The question words that may be added to a graph as question word nodes are underlined and the disjoint answer words that may be added to the graph as answer word nodes are underlined. This illustrates that as "Ireland" is in the question text, it is not included as an answer word for an answer word node in the graph.

[0114] Surface forms and lemmas are converted to nodes of type QuestionWord. For example, "be" is the lemma form of "is", "large" is the lemma form of "largest", and "big" is the lemma form of "biggest".

[0115] Disjoint Question-Answer Pairs:

[0116] The set of pairs generated from this input, along with their frequencies, are then used to build the graph. The following is an example list of the (51) edges between all (17) nodes, from the above example data set:

capital -> largest
capital -> big
capital -> large
capital -> biggest
capital -> population
be -> largest
be -> big
be -> large
be -> biggest
be -> population
dublin -> largest
dublin -> big
dublin -> large
dublin -> biggest
dublin -> population

city -> largest
city -> big
city -> large
city -> biggest
city -> population
is -> largest
is -> big
is -> large
is -> biggest
is -> population
ireland -> largest
ireland -> big
ireland -> large
ireland -> biggest
ireland -> population
be -> africa
be -> cat
be -> family
animal -> largest
animal -> africa
animal -> large
animal -> cat
animal -> family
is -> africa
is -> cat
is -> family
type -> largest
type -> africa
type -> large
type -> cat
type -> family
lion -> largest
lion -> africa
lion -> large
lion -> cat
lion -> family

[0117] Figure 7 shows an extract of an example graph 700 of question and answer vocabulary for the above example. Not all the nodes for the example are shown for purposes of clarity.

[0118] Example QuestionWord nodes are shown for the question words of "capital" 711, "be" 712, "city" 713, "type" 714, "lion" 715 and example AnswerWord nodes are shown for the disjoint answer words of "largest" 721, "biggest" 722, "cat" 723, "family" 724.

[0119] Example edges are shown in the graph 700 as follows:

capital -> largest 731

capital -> biggest 732

be -> largest 733

be -> biggest 734

city -> largest 735

city -> biggest 736

type -> largest 741

type -> family 742

type -> cat 743

lion -> largest 744

lion -> family 745

lion -> cat 746

[0120] Nodes 711-715 and 721-724 and edges 731-746 have weightings based on the frequency of words in the data set and frequency of the disjoint pairs. For example, the weighting of the question word "capital" 711 may be based on a frequency measure of its presence in the question and answer vocabulary. The weighting of edges may be based on a frequency measure of the question-answer pair across the question-answer data. In this simplified example, the edge 733 of "be -> largest" and the edge 734 of "be -> biggest" both appear twice in the question-answer data and will therefore have a higher weighting than if they had appeared once in the question-answer data.

[0121] Unseen Input Questions and Results:

[0122] Sample unseen questions and the n=5 expansions that are produced following graph activation are as follows:

Question 1: "What is the capital of India?"

Expansions: ["big", "biggest", "population", "largest", "large"]

Question 2: "What is a lion?"

Expansions: ["africa", "cat", "family", "largest", "large"]

[0123] The graph 700 may be activated with question nodes corresponding to the question words in the sample unseen questions, such as "capital" 711, "be" 712 (as a lemma of "is") in Question 1 and "lion" 715, "be" 712 (as a lemma form of "is") in Question 2.

[0124] The results of the activation of the graph 700 are the activated nodes of "big", "biggest", "population", "largest", "large" of Question 1 and "africa", "cat", "family", "largest", "large" for Question 2.

[0125] Figure 8 depicts a block diagram of components of the computing device of the query expansion system 110 of Figure 1, in accordance with an embodiment of the present disclosure. It should be appreciated that Figure 8 provides only an illustration of one implementation and does not imply any limitations with regard to the environments in which different embodiments may be implemented. Many modifications to the depicted environment may be made.

[0126] Computing device can include one or more processors 802, one or more computer-readable RAMs 804, one or more computer-readable ROMs 806, one or more computer readable storage media 808, device drivers 812, read/write drive or interface 814, and network adapter or interface 816, all interconnected over a communications fabric 818. Communications fabric 818 can be implemented with any architecture designed for passing data and/or control information between processors (such as microprocessors, communications and network processors, etc.), system memory, peripheral devices, and any other hardware components within the system.

[0127] One or more operating systems 810, and application programs 811, are stored on one or more of the computer readable storage media 808 for execution by one or more of the processors 802 via one or more of the respective RAMs 804 (which typically include cache memory). In the illustrated embodiment, each of the computer readable storage media 808 can be a magnetic disk storage device of an internal hard drive, CD-ROM, DVD, memory stick, magnetic tape, magnetic disk, optical disk, a semiconductor storage device such as RAM, ROM, EPROM, flash memory, or any other computer readable storage media that can store a computer program and digital information, in accordance with embodiments of the disclosure.

[0128] Computing device can also include a R/W drive or interface 814 to read from and write to one or more portable computer readable storage media 826. Application programs 811 on computing device can be stored on one or more of the portable computer readable storage media 826, read via the respective R/W drive or interface 814 and loaded into the respective computer readable storage media 808.

[0129] Computing device can also include a network adapter or interface 816, such as a TCP/IP adapter card or wireless communication adapter. Application programs 811 on computing device can be downloaded to the

computing device from an external computer or external storage device via a network (for example, the Internet, a local area network or other wide area networks or wireless networks) and network adapter or interface 816. From the network adapter or interface 816, the programs may be loaded into the computer readable storage media 808. The network may comprise copper wires, optical fibers, wireless transmission, routers, firewalls, switches, gateway computers and edge servers.

[0130] Computing device can also include a display screen 820, a keyboard or keypad 822, and a computer mouse or touchpad 824. Device drivers 812 interface to display screen 820 for imaging, to keyboard or keypad 822, to computer mouse or touchpad 824, and/or to display screen 820 for pressure sensing of alphanumeric character entry and user selections. The device drivers 812, R/W drive or interface 814, and network adapter or interface 816 can comprise hardware and software stored in computer readable storage media 808 and/or ROM 806.

[0131] The present invention may be a system, a method, and/or a computer program product at any possible technical detail level of integration. 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.

[0132] 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.

[0133] 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.

[0134] 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, configuration data for integrated circuitry, 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 Smalltalk, C++, or the like, and 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.

[0135] 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.

[0136] 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.

[0137] 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.

[0138] 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 blocks 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.

[0139] Cloud Computing

It is to be understood that although this disclosure includes a detailed description on cloud computing, implementation of the teachings recited herein are not limited to a cloud computing environment. Rather, embodiments of the present invention are capable of being implemented in conjunction with any other type of computing environment now known or later developed.

Cloud computing is a model of service delivery for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, network bandwidth, servers, processing, memory, storage, applications, virtual machines, and services) that can be rapidly provisioned and released with minimal management effort or interaction with a provider of the service. This cloud model may include at least five characteristics, at least three service models, and at least four deployment models.

[0140] Characteristics are as follows:

On-demand self-service: a cloud consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with the service's provider.

Broad network access: capabilities are available over a network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs).

Resource pooling: the provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to demand. There is a sense of location independence in that the consumer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter).

Rapid elasticity: capabilities can be rapidly and elastically provisioned, in some cases automatically, to quickly scale out and rapidly released to quickly scale in. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.

Measured service: cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

[0141] Service Models are as follows:

Software as a Service (SaaS): the capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

Platform as a Service (PaaS): the capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.

Infrastructure as a Service (IaaS): the capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

[0142] Deployment Models are as follows:

Private cloud: the cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on-premises or off-premises.

Community cloud: the cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Public cloud: the cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Hybrid cloud: the cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

A cloud computing environment is service oriented with a focus on statelessness, low coupling, modularity, and semantic interoperability. At the heart of cloud computing is an infrastructure that includes a network of interconnected nodes.

Referring now to Figure 9, illustrative cloud computing environment 50 is depicted. As shown, cloud computing environment 50 includes one or more cloud computing nodes 10 with which local computing devices used by cloud consumers, such as, for example, personal digital assistant (PDA) or cellular telephone 54A, desktop computer 54B, laptop computer 54C, and/or automobile computer system 54N may communicate. Nodes 10 may communicate with one another. They may be grouped (not shown) physically or virtually, in one or more networks, such as Private, Community, Public, or Hybrid clouds as described hereinabove, or a combination thereof. This allows cloud computing environment 50 to offer infrastructure, platforms and/or software as services for which a cloud consumer does not need to maintain resources on a local computing device. It is understood that the types of computing devices 54A-N shown in Fig. 9 are intended to be illustrative only and that computing nodes 10 and cloud computing environment 50 can communicate with any type of computerized device over any type of network and/or network addressable connection (e.g., using a web browser).

[0143] Referring now to Figure 10, a set of functional abstraction layers provided by cloud computing environment 50 (Figure 9) is shown. It should be understood in advance that the components, layers, and functions shown in Figure 10 are intended to be illustrative only and embodiments of the invention are not limited thereto. As depicted, the following layers and corresponding functions are provided:

[0144] Hardware and software layer 60 includes hardware and software components. Examples of hardware components include: mainframes 61; RISC (Reduced Instruction Set Computer) architecture based servers 62; servers 63; blade servers 64; storage devices 65; and networks and networking components 66. In some embodiments, software components include network application server software 67 and database software 68.

[0145] Virtualization layer 70 provides an abstraction layer from which the following examples of virtual entities may be provided: virtual servers 71; virtual storage 72; virtual networks 73, including virtual private networks; virtual applications and operating systems 74; and virtual clients 75.

[0146] In one example, management layer 80 may provide the functions described below. Resource provisioning 81 provides dynamic procurement of computing resources and other resources that are utilized to perform tasks within the cloud computing environment. Metering and Pricing 82 provide cost tracking as resources

are utilized within the cloud computing environment, and billing or invoicing for consumption of these resources. In one example, these resources may include application software licenses. Security provides identity verification for cloud consumers and tasks, as well as protection for data and other resources. User portal 83 provides access to the cloud computing environment for consumers and system administrators. Service level management 84 provides cloud computing resource allocation and management such that required service levels are met. Service Level Agreement (SLA) planning and fulfillment 85 provide pre-arrangement for, and procurement of, cloud computing resources for which a future requirement is anticipated in accordance with an SLA.

[0147] Workloads layer 90 provides examples of functionality for which the cloud computing environment may be utilized. Examples of workloads and functions which may be provided from this layer include: mapping and navigation 91; software development and lifecycle management 92; virtual classroom education delivery 93; data analytics processing 94; transaction processing 95; and query expansion processing 96.

[0148] The descriptions of the various embodiments of the present invention have been presented for purposes of illustration, but are not intended to be exhaustive or limited to the embodiments disclosed. Many modifications and variations will be apparent to those of ordinary skill in the art without departing from the scope and spirit of the described embodiments. The terminology used herein was chosen to best explain the principles of the embodiments, the practical application or technical improvement over technologies found in the marketplace, or to enable others of ordinary skill in the art to understand the embodiments disclosed herein.

[0149] Improvements and modifications can be made to the foregoing without departing from the scope of the present invention.

CLAIMS

1. A computer-implemented method for query expansion using a graph of question and answer vocabulary, the method comprising:
 - providing a graph of question word nodes and answer word nodes generated from a set of training data for a given domain in the form of training question texts and training answer texts, wherein the answer word nodes are of disjoint words that do not occur in both a training question text and an associated training answer text, and edges are provided between a disjoint pair of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer;
 - providing weightings for the question word nodes, the answer word nodes, and edges based on frequency data of words in the set of training data;
 - receiving a user query input;
 - activating input nodes in the graph for at least some words in the user query input; and
 - applying spreading activation through the graph using the weightings to result in a top n most highly activated nodes that are output as candidate words for expansion of the user query input.
2. The method as claimed in claim 1, wherein applying spreading activation, propagates a signal in all directions across question word nodes and answer word nodes that are directly or indirectly connected to the input nodes.
3. The method as claimed in claim 1, further comprising capturing user feedback on results of a user query input into a search engine using the candidate words for expansion and using the user feedback to update the graph.
4. The method as claimed in claim 1, wherein a user query input is a question input and activating input nodes activates question word nodes.
5. The method as claimed in claim 4, wherein a user query input includes an answer input in addition to the question input and activating input nodes activates answer word nodes in addition to question word nodes.
6. The method as claimed in claim 5, wherein a user query input includes an answer input in addition to a question input to refine search results to a style of answer and the candidate words are used for refinement of the search results to an answer style.
7. The method as claimed in claim 1, wherein activating input nodes in the graph activates nodes with overlapping surface forms and lemma forms of the words in the user query input.

8. The method as claimed in claim 1, wherein activating input nodes in the graph applies edge relationship definitions for edges between question word and answer word pairs with different weightings for each relationship definition of an edge.
9. The method as claimed in claim 1, wherein providing the graph comprises:
 - providing the set of training data in a given domain in the form of training question texts and training answer texts;
 - identifying disjoint answer words in the training answer text that do not occur in the associated training question text;
 - generating the graph by adding nodes to the graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint words in the set of training data; and
 - adding edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer.
10. A computer-implemented method for query expansion using a graph of question and answer vocabulary, the method comprising:
 - providing a set of training data in a given domain in the form of training question texts and training answer texts;
 - identifying disjoint answer words in a training answer text that do not occur in an associated training question text;
 - adding nodes to a graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint answer words in the set of training data;
 - adding edges to the graph between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer; and
 - providing weightings for the nodes and edges based on frequency data of words in the set of training data.
11. The method as claimed in claim 10, wherein providing weightings for the nodes comprises weighting the nodes by the frequency of occurrence of the question word of a question word node in the set of training data and the frequency of occurrence of the answer word of an answer word node in the set of training data.
12. The method as claimed in claim 10, wherein providing weightings of edges between nodes comprises weighting an edge according to the frequency of the disjoint question word and answer word pair across the set of training data.

13. The method as claimed in claim 10, wherein adding edges to the graph comprises creating edges only between question word and answer word pairs with the question words and answer words having lexical or semantic relationships.
14. The method as claimed in claim 10, wherein providing the graph includes providing relationship definitions for edges between question word and answer word pairs and providing a different weighting for at least one relationship definition of one edge.
15. The method as claimed in claim 10, wherein adding nodes to the graph as question word nodes comprises conditionally selecting words from the training question texts according to predefined constraints.
16. A system for query expansion using a graph of question and answer vocabulary, the system comprising:
a processor; and
a memory communicatively coupled to the processor and configured to store computer program instructions;
wherein the processor is configured to execute the computer program instructions to:
provide a graph of question word nodes and answer word nodes generated from a set of training data for a given domain in the form of training question texts and training answer texts, wherein the answer word nodes are of disjoint words that do not occur in both a training question text and an associated training answer text, and edges are provided between a disjoint pair of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer;
provide weightings for the nodes and edges based on frequency data of words in the set of training data;
receive a user query input;
activate input nodes in the graph for at least some words in the user query input;
apply spreading activation through the graph using the weightings; and
return a result of a top n most highly activated nodes that are output as candidate words for expansion of the user query input.
17. The system as claimed in claim 16, wherein the processor is further configured to capture user feedback on results of a user query input into a search engine using the candidate words for expansion and to update the graph using the user feedback.
18. The system as claimed in claim 16, wherein the user query input is a question input, and/or an answer input and the graph input component activates question word nodes and/or answer word nodes.

19. The system as claimed in claim 18, wherein a user query input includes an answer input in addition to the question input to refine search results to a style of answer, and wherein the graph output component returns a result of a top n most highly activated nodes that are used as candidate words for refinement of the search results to an answer style.
20. A system for query expansion using a graph of question and answer vocabulary, the system comprising:
a processor; and
a memory communicatively coupled to the processor and configured to store computer program instructions;
wherein the processor is configured to execute the computer program instructions to:
provide a set of training data in a given domain in the form of training question texts and training answer texts;
identify disjoint answer words in a training answer text that do not occur in both an associated training question text;
add nodes to a graph as question word nodes and answer word nodes, wherein the answer word nodes are for the identified disjoint answer words in the set of training data;
add edges between a disjoint pair of a question word in a training question and a disjoint answer word in an associated training answer; and
provide weightings for the nodes and edges based on frequency data of the words in the set of training data.
21. The system as claimed in claim 20, wherein the processor is further configured to:
provide weightings for the nodes configured to weight the nodes by the frequency of occurrence of the question word of a question word node in the set of training data and the frequency of occurrence of the answer word of an answer word node in the set of training data; and
provide weightings of edges between nodes configured to weight an edge according to the frequency of the disjoint question word and answer word pair across the set of training data.
22. The system as claimed in claim 20, wherein the processor is configured to create edges only between question word and answer word pairs with the question words and answer words having lexical or semantic relationships.
23. The system as claimed in claim 20, wherein the processor is configured to provide relationship definitions for edges between question word and answer word pairs and providing a different weighting for at least one relationship definition of one edge.

24. The system as claimed in claim 20, wherein the processor is configured to receive user feedback on results of a user query input into a search engine using the candidate words for expansion and a graph updating component for applying the user feedback to the graph.

25. A computer program product for query expansion using a graph of question and answer vocabulary, the computer program product comprising a computer readable storage medium having program instructions embodied therewith, the program instructions executable by a processor to cause the processor to:

- provide a graph of question word nodes and answer word nodes generated from a set of training data for a given domain in the form of training question texts and training answer texts, wherein the answer word nodes are of disjoint words that do not occur in both a training question text and an associated training answer text, and edges are provided between a disjoint pair of a question word node for a question word in a training question and an answer word node for a disjoint answer word in an associated training answer;


- provide weightings for the nodes and edges based on frequency data of the words in the set of training data;

- receive a user query input;

- activate input nodes in the graph for at least some words in the user query input; and

- apply spreading activation through the graph using the weightings to result in a top n most highly activated nodes that are used as candidate words for expansion of the user query input.

1/10

**FIG. 1**

2/10

FIG. 2A




FIG. 2B







FIG. 2C




3/10

300
↓**FIG. 3**


4/10

400
↓**FIG. 4**

5/10

**FIG. 5**

6/10

**FIG. 6**

7/10






FIG. 7

8/10

**FIG. 8**

9/10

**FIG. 9**

10/10




FIG. 10

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB2019/055380

A. CLASSIFICATION OF SUBJECT MATTER

G06F 17/27(2006.01)i

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

VEN; DWPI; CNABS; CNKI: query expansion, word, graph, question, answer, node, vocabulary, training, learning

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	CN 107832319 A (UNIV BEIJING TECHNOLOGY) 23 March 2018 (2018-03-23) see the whole document	1-25
A	US 2009070311 A1 (AT & T CORP) 12 March 2009 (2009-03-12) see the whole document	1-25
A	CN 105659225 A (MICROSOFT TECHNOLOGY LICENSING LLC) 08 June 2016 (2016-06-08) see the whole document	1-25
A	US 2016147872 A1 (IBM) 26 May 2016 (2016-05-26) see the whole document	1-25



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

“A” document defining the general state of the art which is not considered to be of particular relevance

“E” earlier application or patent but published on or after the international filing date

“L” document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

“O” document referring to an oral disclosure, use, exhibition or other means

“P” document published prior to the international filing date but later than the priority date claimed

“T” later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

“X” document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

“Y” document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

“&” document member of the same patent family

Date of the actual completion of the international search

29 September 2019

Date of mailing of the international search report

12 October 2019

Name and mailing address of the ISA/CN

National Intellectual Property Administration, PRC
6, Xitucheng Rd., Jimen Bridge, Haidian District, Beijing
100088
China

Authorized officer

WEN,Rui

Facsimile No. (86-10)62019451

Telephone No. 62411119

INTERNATIONAL SEARCH REPORT
Information on patent family members

International application No.

PCT/IB2019/055380

Patent document cited in search report			Publication date (day/month/year)	Patent family member(s)		Publication date (day/month/year)	
CN	107832319	A	23 March 2018	None			
US	2009070311	A1	12 March 2009	US	8543565 B2	24 September 2013	
CN	105659225	A	08 June 2016	US	9286396 B2	15 March 2016	
				EP	3025254 A1	01 June 2016	
				US	2015032767 A1	29 January 2015	
				WO	2015013469 A1	29 January 2015	
US	2016147872	A1	26 May 2016	US	2016147870 A1	26 May 2016	
				US	10198504 B2	05 February 2019	
				US	10055485 B2	21 August 2018	