



- (51) **International Patent Classification:**
G06F 17/30 (2006.01)
- (21) **International Application Number:**
PCT/US2014/024271
- (22) **International Filing Date:**
12 March 2014 (12.03.2014)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/842,572 15 March 2013 (15.03.2013) US
- (71) **Applicant:** GOOGLE INC. [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).
- (72) **Inventors:** MOXLEY, Emily; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US). LIU, Sean; 1600 Amphitheatre Parkway, Mountain View, CA 94903 (US).
- (74) **Agent:** CHAPMAN, Ernest, F.; Finnegan, Henderson, Farabow, Garrett &, Dunner, LLP, 901 New York Avenue, NW, Washington, DC 20001-4113 (US).
- (81) **Designated States** (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM,

AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, 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, 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))

(54) **Title:** PROVIDING SEARCH RESULTS USING AUGMENTED SEARCH QUERIES




FIG. 4

(57) **Abstract:** Methods and systems are provided for updating a set of results. In some implementations, a search query associated with an entity reference is received. The entity reference corresponds to one or more distinct entities. A set of results for the search query is provided, and the set of results distinguishes between distinct entities. One or more attributes for at least one entity of the one or more distinct entities are identified based at least in part on the set of results. One or more additional search queries are identified based on the search query, the at least one entity, and the one or more attributes. An input selecting at least one of the additional search queries is received. An updated set of results is provided based on the selected additional search queries. The updated set of results comprises at least one result not in the set of results.

PROVIDING SEARCH RESULTS USING AUGMENTED SEARCH QUERIES

Background

[0001] This disclosure relates generally to providing search results in response to a search query containing an entity reference. Search engines receive search queries containing a reference to a person, such as the person's name. Results to these queries are often times not sufficiently organized, not comprehensive enough, or otherwise not presented in a useful way.

Summary

[0002] In some implementations, a system receives a search query containing an entity reference, such as a person's name, that corresponds to one or more distinct entities. The system provides a set of results, where each result is associated with at least one of the distinct entities. The system uses the set of results to identify attributes of the entity and uses the identified attributes to generate additional, augmented search queries associated with the entity. The system updates the set of results based on one or more of these augmented search queries.

[0003] In some implementations, a computer-implemented method includes receiving a search query associated with an entity reference, wherein the entity reference corresponds to one or more distinct entities. The method includes providing a set of results for the search query where the set of results distinguishes between distinct entities. The method includes identifying one or more attributes of at least one entity of the one or more distinct entities based at least in part on the set of results. The method includes generating one or more additional search queries based on the search query, the at least one entity, and the one or more attributes. The method includes receiving an input selecting at least one of the one or more additional search queries and providing an updated set of results based on the selected one or more additional search queries, where the updated set of results comprises at least one result not in the set of results.

[0004] These and other implementations can each include one or more of the following features. In some implementations, generating one or more additional search queries comprises ranking the identified one or more attributes and generating one or more additional search queries based on the search query, the at least one entity, the one or more

attributes, and the ranking. In some implementations, the ranking is based on frequency of occurrence. In some implementations, the ranking is based on a location of each of the one or more attributes with respect to the at least one entity in the set of results. In some implementations, the method further comprises providing to a user the one or more additional search queries. In some implementations, the receiving an input comprises receiving an input from the user.

[0005] In some implementations, a system includes one or more computers configured to perform operations. The operations include receiving a search query associated with an entity reference, wherein the entity reference corresponds to one or more distinct entities. The operations include providing a set of results for the search query where the set of results distinguishes between distinct entities. The operations include identifying one or more attributes of at least one entity of the one or more distinct entities based at least in part on the set of results. The operations include generating one or more additional search queries based on the search query, the at least one entity, and the one or more attributes.

[0006] In some implementations, a computer-implemented method includes receiving a search query associated with an entity reference, wherein the entity reference corresponds to one or more distinct entities. The method includes providing a set of results for the search query where the set of results distinguishes between distinct entities. The method includes identifying one or more attributes of at least one entity of the one or more distinct entities based at least in part on the set of results. The method includes generating one or more additional search queries based on the search query, the at least one entity, and the one or more attributes. The method includes receiving an input selecting at least one of the one or more additional search queries and providing an updated set of results based on the selected one or more additional search queries, where the updated set of results comprises at least one result not in the set of results.

Brief Description of the Figures

[0007] FIG. 1 is a high level block diagram of a system for providing search results in accordance with some implementations of the present disclosure;

[0008] FIG. 2 shows an illustrative example of a user interface for providing a set of search results in accordance with some implementations of the present disclosure;

[0009] FIG. 3 shows an illustrative natural language processing diagram in accordance with some implementations of the present disclosure;

[0010] FIG. 4 shows an illustrative example of a user interface for selecting one or more augmented search queries in accordance with some implementations of the present disclosure;

[0011] FIG. 5 shows an illustrative example of a user interface for updating search results in accordance with some implementations of the present disclosure;

[0012] FIG. 6 is an exemplary diagram of augmented search queries in accordance with some implementations of the present disclosure;

[0013] FIG. 7 shows an illustrative knowledge graph portion containing nodes and edges in accordance with some implementations of the present disclosure;

[0014] FIG. 8 shows an illustrative knowledge graph portion in accordance with some implementations of the present disclosure;

[0015] FIG. 9 shows an illustrative knowledge graph portion in accordance with some implementations of the present disclosure;

[0016] FIG. 10 shows a flow diagram of illustrative steps for providing search results in accordance with some implementations of the present disclosure;

[0017] FIG. 11 shows an illustrative system in accordance with some implementations of the present disclosure; and

[0018] FIG. 12 is a block diagram of an illustrative computer system in accordance with some implementations of the present disclosure.

Detailed Description

[0019] FIG. 1 is a high level block diagram of a system for providing search results in accordance with some implementations of the present disclosure. In some implementations, a system 100 receives a search query 102 associated with an entity reference. Processing block 106 generates additional search queries, such as augmented search queries, based on an entity corresponding to the entity reference, and uses the additional search queries to provide a set of results 110 to the search query. Search queries may be textual, image-based, audio-based, video-based, based in any other suitable format, or any combination thereof. For example, system 100 may receive a textual search query in a query box. For purposes of brevity, system 100 will also be referred to herein as the system.

[0020] Entity reference, as used herein, refers to an identifier that corresponds to one or more distinct entities. Entity, as used herein, refers to a thing or concept that is singular, unique, well defined, and distinguishable. An entity reference may correspond to more than one distinct entity. For example, an entity reference may be a person's name, and

corresponding entities may include distinct people who share the referenced name. For purposes of brevity and clarity, and not by way of limitation, the present disclosure is presented with respect to the entities being people and the entity references being names of or other references to people. It will be understood that any suitable type of entity, or combinations thereof, may be used in the context of the present disclosure, including, for example, a person, place, item, idea, topic, abstract concept, concrete element, other suitable thing, or any combination thereof. In an example, an entity may be the physical person John Adams, and an entity reference may be an identified that refers to John Adams, such as his name.

[0021] The system provides search results in response to a search query. Search results may be associated with an entity corresponding to an entity reference in the search query. The system may provide a set of results that distinguishes between distinct entities that correspond to the entity reference. In an example, distinct entities may include three people who are associated with the name John Adams, and the entity reference may be the name John Adams. The system may parse each result in the set of results to identify three different people associated with the name John Adams: John Adams the Second President, John Quincy Adams the Sixth President, and John Adams the artist. This parsing may be accomplished using any suitable technique.

[0022] In some implementations, the system uses natural language processing to parse a search query. As used herein, natural language refers to words, syntax, and other language such as it could be used in conversation or prose. For example, natural language may include complete sentences, questions, idiom, punctuation, any other suitable language elements or structures, or any combination thereof. For example, the question [Who was the first person to fly an airplane?] is a natural language question. In contrast, formal language follow relatively more constrained rules of grammar and syntax. An example of formal language is a computer programming language such as C or BASIC. It will be understood that queries, including natural language queries, may be in any suitable language such as English, French, Chinese, and so on. It will be understood that in some implementations, the system need not receive a natural language query and may receive a query in any suitable form. It will also be understood that the system may receive questions, provide questions for answering, receive responses, provide answers, and perform any other suitable steps using natural language, formal language, keywords, voice, video, images, any other suitable communication technique, or any combination thereof. Natural language processing is described in detail below in relation to **FIG. 3**.

[0023] The system identifies one or more attributes of at least one entity, where the at least one entity corresponds to the entity reference. Attribute, as used herein, refers to any feature or characteristic associated with an entity that the system may identify based on the set of results. In an example, for the entity John Adams, the system might identify, in the set of results, attributes including [second president], [Abigail Adams], and [Alien and Sedition Acts]. The system generates one or more additional queries, such as one or more augmented versions of the original search query, for an entity corresponding to the identified entity reference using the identified attributes.

[0024] Referring back to FIG. 1, system 100 includes search query block 102, structure data block 104, processing block 106, content block 108, and output block 110. System 100 may include any suitable hardware, software, or both for implementing the features described in the present disclosure and will generally be referred to, herein, as the system. In some implementations, system 100 receives a search query containing an entity reference in search query block 102. Processing block 106 processes the search query from search query block 102, data from structured data block 104, and content from content block 108 to determine a set of results that is provided to output block 110. In some implementations, processing block 106 identifies one or more entities corresponding to identical to similar entity references from the set of results and distinguishes between the distinct, identified entities. Processing block 106 identifies one or more attributes of an identified entity based on the search query from search query block 102, information from structured data block 104 and content block 108, and the set of results. Processing block 106 generates additional search queries based on the search query from search query block 102 and the identified attributes. The additional search queries may be displayed in output block 110. In some implementations, output block 110 is used to update a set of results based on the additional search queries generated in processing block 106.

[0025] Structured data block 104 may include a data structure containing structured or otherwise organized, information. In some implementations, structured data block 104 contains one or more of, data graphs, databases, indexes, any other suitable collection of content, or any combination thereof.

[0026] In some implementations, structured data block 104 includes a data graph. In some implementations, a data graph includes data organized in a graph containing nodes and edges. The data of a data graph may include information about relationships between things and concepts, and this information may be represented as nodes and edges of a graph. The nodes of a data graph each contain a piece or pieces of data and the edges represent

relationships between the data contained in the nodes that the edges connect. A knowledge graph, which is a particular implementation of a data graph, is described below in FIGS. 6–8.

[0027] Content block 108 includes data related to web sites and other such searchable and/or indexed data sources. In an example, content block 108 includes an index of webpages on the internet containing text, images, videos, links, other suitable content, and any combination thereof. In some implementations, content block 108 includes information from a database, private intranet, public network, private network, any other suitable collection of information, or any combination thereof. In some implementations, content block 108 includes content used by a search engine such as a query processing engine to determine search results. It will be understood that content block 108 may include both structured and unstructured data. Content block 108 may include data derived from unstructured data. For example, content block 108 may include structured data culled from websites, i.e., unstructured data, regarding entity references and associated attributes as determined from the source websites.

[0028] System 100 provides a set of results in processing block 106 for the search query from search query block 102. System 100 may determine the set of results based on data from data block 104, data from content block 108, any other suitable data, or any combination thereof. In some implementations, processing block 106 identifies entities in the set of results that correspond to the entity reference received in search query block 102.

[0029] System 100 identifies one or more attributes for at least one entity in processing block 106. In some implementations, system 100 may select the at least one entity for which attributes will be identified in processing block 106 based on content block 106, structured data block 104, any other suitable data, or any combination thereof.

[0030] System 100 generates one or more additional search queries in processing block 106 based on the identified attributes. In some implementations, processing block 106 ranks the identified attributes and generates one or more additional search queries based on the identified attributes and the ranking. Generating additional search queries is discussed in detail below in relation to FIG. 5. In some implementations, processing block 106 provides the additional search queries to output block 110 for display.

[0031] System 100 provides an updated set of results in output block 110. System 100 may update the set of results based on one or more user-selected additional search queries in processing block 106. In some implementations, system 100 bases the selection on data from data block 104, content from content block 108, any other suitable content, or any combination thereof.

[0032] FIG. 2 shows an illustrative example of a user interface for providing a set of results in accordance with some implementations of the present disclosure. In some implementations, user interface 200 includes search query box 218 containing received search query 202. In some implementations, search query 202 is received using textual input, voice command, image drag-and-drop, gesture recognition, camera input, any other suitable technique, or any combination thereof. In some implementations, search button 220 receives user input indicating a desired search. The search button may be activated, for example, using input received using a mouse or touchpad, and activation of the search button causes search results to be presented. Search query 202 and corresponding search results, including, for example, SearchResult1 204, SearchResult2 206, and SearchResult3 208 are shown. The system may parse search query 202 to identify one or more associated entity references. In some implementations, the system uses natural language processing to parse a search query 202. Natural language processing is discussed in detail below in relation to FIG. 3.

[0033] In some implementations, the system identifies Entity1 212, Entity2 214, and Entity3 216 from the search results based on whether the entity reference corresponds to the entity. In some implementations, the system identifies entities 212, 214, and 216 by parsing each result in the set of results 204, 206, and 208. In some implementations, the system uses natural language processing to parse each result in the set of results. In some implementations, each result in the set of results is associated respectively with at least one entity. In some implementations, the system displays identified entities 212, 214, and 216 in entity selection panel 210. In another example, identified entities 212, 214, and 216 may be displayed in a profile card or other suitable medium for display. Entity selection box 210 may include a mechanism for selection and de-selection of entities 212, 214, and 216. In the illustrated example, entity selection panel 210 provides clickable hyperlinks for entities 212, 214, and 216. that may receive user input for selection and de-selection of the entities. In another example, checkboxes may be used for selection and de-selection of entities 212, 214, and 216. In some implementations, the system may select one or more entities, for example Entity2 214.

[0034] FIG. 3 shows illustrative natural language processing diagram 300 in accordance with some implementations of the present disclosure. In some implementations, the search system uses natural language processing to parse a search query. Parsing may include identifying words, parts of speech, meanings, and relationships between words in the query.

[0035] Diagram 300 shows the natural language processing of the search query "Tom Hanks Movies in 2010." In some implementations, the search system divides the query into search units "Tom Hanks" 304, "Movies" 302, "In" 306 and "2010" 308. In some implementations, the search system divides the search query based on a phrase dictionary, search history, user preferences, predetermined parameters, system settings, any other suitable parameters, or any combination thereof. In some implementations, the search system determines the part of speech of each search unit. For example, the search system may determine that "Movies" 302 is a common noun, "Tom Hanks" 304 is a proper noun, "In" 306 is a preposition, and "2010" 308 is a number. It will be understood that this particular part of speech notation is merely an example and that any suitable identification, notation, or both may be used. In some implementations, the search system determines the semantic function or dependency of the search unit in the search query. For example, the search system may determine that "Movies" 302 is the root of the query, "Tom Hanks" 304 is a noun modifier, which means that it modifies the root, "In" 306 is a preposition, and "2010" 308 is the object of the preposition. In some implementations, the search system assigns a directional relationship between the search units as illustrated in diagram 300. For example, "In" 306 and "Tom Hanks" 304 is a child node of the root "Movies" 302, and 2010 is a child of "In" 306. In some implementations, natural language processing is used to identify related nodes in a knowledge graph, uses information from the knowledge graph, relates to a data structure such as the knowledge graph in any suitable way, or any combination thereof. It will also be understood that the search system may perform natural language processing without relying on the knowledge graph.

[0036] FIG. 4 shows an illustrative example of a user interface 400 for providing additional search queries in accordance with some implementations of the present disclosure. In some implementations, user interface 400 includes search query box 412 and search button 416. Search query 402 is shown and corresponding identified entities are presented in entity selection panel 404. In some implementations, the system may provide search results associated with at least one selected entity. For example, Entity2 414 is shown as selected in entity selection panel 404. Search results for Entity2 414, are shown, for example, as search results 406. In some implementations, the system provides entity panel 404 for switching between entity selections and for selecting more than one entity.

[0037] The system identifies one or more attributes for selected Entity2 414. In some implementations, the system parses each result in a set of results associated with the selected entity to identify the one or more attributes. In some implementations the system uses natural

language processing, as described above in relation to FIG. 3, to parse each result in a set of results. For example, in a web search engine context, the system may be configured to access and parse webpages in order to identify entity references contained in the webpages as well as attribute information associated with respective entity references also contained in the websites. Attribute information can also be derived from structured data, such as structured data 104, in which entity references are looked up and associated attribute information can be readily identified.

[0038] In some implementations, the system generates one or more additional search queries based on the search query, the at least one selected entity, and the one or more attributes. In some implementations, the system may augment the search query with the one or more attributes to generate one or more additional queries 418. In some implementations, the system ranks the one or more attributes and generates one or more additional search queries based on the search query, the at least one entity, the one or more attributes, and the ranking of the one or more attributes. In some implementations, the system may use the ranking to order the display of the one or more additional search queries. For example, the system may display a list of additional search queries, and in the list, an additional search query based on a higher ranking attribute may appear nearer the top of the list than an additional search query based on a lower ranking attribute. In some implementations the system may use the ranking to determine which of the one or more attributes to use in generating one or more additional search queries. For example, if an attribute has a particularly low ranking relative to the other of the one or more attributes, the system may determine not to use the attribute to generate an additional search query.

[0039] In some implementations, the ranking is based on the number of occurrences, frequency of occurrence, or both of one or more attributes in a document, such as a webpage, where number is an absolute count within a document and frequency is a count within the document normalized by the length of the document. In some implementations, the system determines the ranking based in part on the number of occurrence, frequency of occurrence, or both, of the one or more attributes across a number of documents. In an example, an attribute occurring at a location proximate to the selected entity may be ranked higher than an attribute occurring at a location distal to the selected entity. The location with respect to the layout of the document may be based on, for example, whether an attribute occurs at the top of the page, in the first sentence of a paragraph, within a title, within a picture caption, or in any other suitably prominent place in the document. In some implementations, the system

determines the ranking based in part on one or more of the location, number of occurrence, and frequency of occurrence.

[0040] In some implementations, the system determines a ranking based in part on the document where the attribute occurs. For example, a webpage may be associated with a popularity score, a freshness score, a rating based on the number of hyperlinks to and from that page, a manual ranking, any other suitable metric, or any combination thereof. In some implementations, the system determines the ranking based in part on one or more of those document rankings.

[0041] In some implementations, the system may scale, normalize, weigh, combine with other data, or otherwise adjust the ranking based on page quality, freshness, popularity, user input, system design, any other suitable criteria, or any combination thereof. In an example, the rank of a first attribute from a recently updated webpage may be weighted with a higher weight than the rank of a second attribute from an older webpage. In another example, the ranking across a number of webpages may be normalized such that each document has the same relative contribution to an aggregate ranking. In another example, ranks from highly visited webpages are assigned a higher weight than ranks from infrequently visited websites. It will be understood that any of the foregoing ranking techniques, any other suitable ranking technique, or combination thereof, may be used to determine a ranking of the one or more attributes.

In some implementations, the system displays the one or more additional search queries, for example, on the search screen in query selection box 418.

[0042] FIG. 5 shows an illustrative example of a user interface 500 for updating search results in accordance with some implementations of the present disclosure. In some implementations, user interface 500 includes search query box 512 and search button 516. One or more additional search queries, for example, [Entity2+Attribute1] 518, [Entity2+Attribute2] 520, [Entity2+Attribute3] 522, [Entity2+Attribute4] 524, are shown in query selection box 504. In some implementations, the system provides a mechanism for selection and de-selection of the one or more additional search queries. For example, checkboxes are an exemplary mechanism for selection and de-selection of the one or more additional search queries. [Entity2+Attribute1] 518 is shown as selected in the illustrated example. In some implementations, the system provides an updated set of search results using the at least one selected additional search query. More particularly, the system provides search results for the search query represented by selected query 518 in place of the original search query 502. This effectively allows a user to perform a more comprehensive

search on an entity having a specific attribute that may not have been explicit in the original search query 502, and therefore may not have resulted in as many useful results when searched based on original search query 502.

[0043] FIG. 6 is exemplary diagram 600 of illustrative additional search queries in accordance with some implementations of the present disclosure. The illustrated diagram shows how the system can process an original search query for [Kate Middleton] in order to generate and augmented search queries [Kate Middleton] + [wedding], [Kate Middleton] + [St. Andrews University], and [Kate Middleton] + [sister]. Additional search queries 604, 606, and 608 are shown with respective corresponding circles 612, 614, and 616 corresponding to respective sets of results are shown. In the illustrated example, a set of results for selected entity [Kate Middleton] 602 is shown in as an ellipse with a dashed line. The system identifies three attributes of [Kate Middleton] from the set of results for the selected entity [Kate Middleton] 602. In the example, three identified attributes—[St. Andrews University], [wedding], and [sister]—are shown. The system generates an additional search query for each of the identified attributes. In the illustrated example, additional search queries [Kate Middleton] + [St. Andrews University] 608, [Kate Middleton] + [sister] 606, and [Kate Middleton] + [wedding] 604 are shown. The system may provide these three additional search queries to a user for selection of one or more to be used to update the set of results. For example, updating the set of results based on additional search query [Kate Middleton] + [sister] 606 would result in the updated set of results represented by the corresponding checkered circle 606. Generally, it will be understood that each of the one or more additional search queries may result in updated sets of results that overlap with one or more of the other additional search queries. For example, a webpage discussing Kate Middleton's wedding and her sister Pippa Middleton as the maid of honor may appear in the updated set of results for additional query [Kate Middleton] + [sister] 606 and in the updated set of results for additional query [Kate Middleton] + [wedding] 604. It will also be understood that the updated set of results for each of the one or more additional queries includes at least one reference that is not in the set of results for the at least one entity. As shown in the illustrated example, the circle with the dashed line corresponding to the set of results associated with the entity [Kate Middleton] 602 includes some but not all results found in each of the three updated set of results 612, 614, and 616 for additional search queries 604, 606, and 608. In an example, the additional search query [Kate Middleton] + [sister] would return the set of results shown in the figure as a checkered circle 606, which

includes webpages that are within and without the set of results 610 for [Kate Middleton] 602.

[0044] FIG. 7 shows an illustrative data graph containing nodes and edges in accordance with some implementations of the present disclosure. In some implementations, data in data structure 104 of FIG. 1 is stored as a data graph. In some implementations, illustrative data graph 700 is a portion of a knowledge graph. The knowledge graph is a particular implementation of a data graph and will be described in further detail in relation to FIGS. 8 and 9 below. It will be understood that the data graph implementation of FIG. 6, and the knowledge graph, are merely examples of a data structure that may be used by the system to store entity references and other data, and that any suitable data format may be used. For example, data in data structure 104 of FIG. 1 may be stored as a list of entities and associated entity types. Data stored by the data structure may include any suitable data such as references to data, text, images, characters, computer files, databases, any other suitable data, or any combination thereof. It will be understood that in some implementations, the node and edge description is merely illustrative and that the construction of the data structure may include any suitable technique for describing information and relationships. In an example, nodes may be assigned a unique identification number, and an edge may be described using the identification numbers that a particular edge connects. It will be understood that the representation of data as a graph is merely exemplary and that data may be stored, for example, as a computer file including pieces of data and links and/or references to other pieces of data.

[0045] In some implementations, data may be organized in a database using any one or more data structuring techniques. For example, data may be organized in a graph containing nodes connected by edges. In some implementations, the data may include statements about relationships between things and concepts, and those statements may be represented as nodes and edges of a graph. The nodes each contain a piece or pieces of data and the edges represent relationships between the data contained in the nodes that the edges connect. In some implementations, the graph includes one or more pairs of nodes connected by an edge. The edge, and thus the graph, may be directed, *i.e.* unidirectional, undirected, *i.e.* bidirectional, or both, *i.e.* one or more edges may be undirected and one or more edges may be directional in the same graph. Nodes may include any suitable data or data representation. Edges may describe any suitable relationships between the data. In some implementations, an edge is labeled or annotated, such that it includes both the connection between the nodes, and descriptive information about that connection. A particular node may be connected by

distinct edges to one or more other nodes, or to itself, such that an extended graph is formed. For purposes of clarity, a graph based on the structure described immediately above is referred to herein as a knowledge graph. In some implementations, the knowledge graph may be a useful for representing information and in providing information in search.

[0046] FIG. 7 shows an illustrative data graph containing nodes and edges in accordance with some implementations of the present disclosure. In some implementations, illustrative data graph 700 is a portion of a knowledge graph. The knowledge graph will be described in further detail in relation to FIGS. 8 and 9 below. It will be understood that the data graph implementation of FIG. 7, and in particular the knowledge graph, is merely an example of a data structure that may be used by the system and that any suitable data structure or data format may be used. Data stored by the data structure may include any suitable data such as references to data, text, images, characters, computer files, databases, any other suitable data, or any combination thereof. It will be understood that in some implementations, the node and edge description is merely illustrative and that the construction of the data structure may include any suitable technique for describing information and relationships. In an example, nodes may be assigned a unique identification number, and an edge may be described using the identification numbers that a particular edge connects. It will be understood that the representation of data as a graph is merely exemplary and that data may be stored, for example, as a computer file including pieces of data and links and/or references to other pieces of data.

[0047] In some implementations, data may be organized in a database using any one or more data structuring techniques. For example, data may be organized in a graph containing nodes connected by edges. In some implementations, the data may include statements about relationships between things and concepts, and those statements may be represented as nodes and edges of a graph. The nodes each contain a piece or pieces of data and the edges represent relationships between the data contained in the nodes that the edges connect. In some implementations, the graph includes one or more pairs of nodes connected by an edge. In some implementations, the edge, and thus the graph, may be directed, undirected, or both. In an example, directed edges form a unidirectional connection. In an example, undirected edges form bidirectional connections. In an example, a combination of both directed and undirected edges may be included in the same graph. Nodes may include any suitable data or data representation. Edges may describe any suitable relationships between the data. In some implementations, an edge is labeled or annotated, such that it includes both the connection between the nodes, and descriptive information about that

connection. It will be understood that in some implementations, edges between data sources need not be labeled. A particular node may be connected by distinct edges to one or more other nodes, or to itself, such that an extended graph is formed.

[0048] In some implementations, the grouping of an edge and two nodes is referred to as a triple. The triple represents the relationship between the nodes, or in some implementations, between the node and itself. In some implementations, higher order relationships are modeled, such as quaternary and n -ary relationships, where n is an integer greater than 2. In some implementations, information modeling the relationship is stored in a node, which may be referred to as a mediator node. In an example, the information "Person X Donates Artifact Y To Museum Z" is stored in a mediator node connected entity nodes to X, Y, and Z, where each edge identifies the role of each respective connected entity node.

[0049] Illustrative graph 700 includes nodes 702, 704, 706, and 708. Data graph 700 includes edge 710 connecting node 702 and node 704. Data graph 700 includes edge 712 connecting node 702 and node 706. Data graph 700 includes edge 714 connecting node 704 and node 708. Data graph 700 includes edge 716 and edge 718 connecting node 702 and node 708. Data graph 700 includes edge 720 connecting node 708 to itself. Each aforementioned group of an edge and one or two distinct nodes may be referred to as a triple or 3-tuple. As illustrated, node 702 is directly connected by edges to three other nodes, while nodes 704 and 708 are directly connected by edges to two other nodes. Node 706 is connected by an edge to only one other node, and in some implementations, node 706 is referred to as a terminal node. As illustrated, nodes 702 and 708 are connected by two edges, indicating that the relationship between the nodes is defined by more than one property. As illustrated, node 708 is connected by edge 720 to itself, indicating that a node may relate to itself. While illustrative data graph 700 contains edges that are not labeled as directional, it will be understood that each edge may be unidirectional or bidirectional. It will be understood that this example of a graph is merely an example and that any suitable size or arrangement of nodes and edges may be employed.

[0050] FIG. 8 shows illustrative knowledge graph portion 800 in accordance with some implementations of the present disclosure. A knowledge graph is a particular implementation of a data graph as illustrated above in relation to data graph 700 of FIG. 7.

[0051] Generally, nodes in a knowledge graph can be grouped into several categories. Nodes may represent, for example, entities, organizational data such as entity types and properties, literal values, and models of relationships between other nodes.

[0052] An edge in a knowledge graph may represent a semantic connection defining a relationship between two nodes. The edge may represent a prepositional statement such as [Is A], [Has A], [Is Of A Type], [Has Property], [Has Value], any other suitable statement, or any combination thereof. For example, the entity node of a particular person may be connected by a [Date Of Birth] edge to a terminal node containing a literal of his or her specific date of birth. In some implementations, the properties defined by edge connections of an entity may relate to nodes connected to the type of that entity. For example, the entity type node [Movie] may be connected to entity nodes [Actor] and [Director], and a particular movie may be connected by an edge property [Has Actor] to an entity node representing a particular actor.

[0053] In some implementations, a node of a knowledge graph represents an entity. An entity is a thing or concept that is singular, unique, well-defined and distinguishable. For example, an entity may be a person, place, item, idea, abstract concept, concrete element, other suitable thing, or any combination thereof. It will be understood that in some implementations, the data graph contains an entity reference, and not the physical embodiment of the entity. For example, an entity may be the physical embodiment of George Washington, while an entity reference is an abstract concept that refers to George Washington. In another example, the entity "New York City" refers to the physical city, and the data graph uses a concept of the physical city as represented by, for example, an element in a data structure, the name of the entity, any other suitable element, or any combination thereof. Where appropriate, based on context, it will be understood that the term entity as used herein may correspond to an entity reference, and the term entity reference as used herein may correspond to an entity.

[0054] Generally, entities include things or concepts represented linguistically by nouns. For example, the color [Blue], the city [San Francisco], and the imaginary animal [Unicorn] may each be entities. An entity reference generally refers to the concept of the entity. For example, the entity reference [New York City] refers to the physical city, and the data graph uses a concept of the physical city as represented by, for example, an element in a data structure, the name of the entity, any other suitable element, or any combination thereof.

[0055] In some implementations, a node representing organizational data may be included in a knowledge graph. These may be referred to herein as entity type nodes. As used herein, an entity type node may refer to a node in a knowledge graph, while an entity type may refer to the concept represented by an entity type node. An entity type may be a defining characteristic of an entity. For example, entity type node Y may be connected to an

entity node X by an [Is A] edge or link, discussed further below, such that the graph represents the information "The Entity X Is Type Y." For example, the entity node [George Washington] may be connected to the entity type node [President]. An entity node may be connected to multiple entity type nodes, for example, [George Washington] may also be connected to entity type node [Person] and to entity type node [Military Commander]. In another example, the entity type node [City] may be connected to entity nodes [New York City] and [San Francisco]. In another example, the concept [Tall People], although incompletely defined as a definition of the property [tall] is not necessarily known, may exist as an entity type node. In some implementations, the presence of the entity type node [Tall People], and other entity type nodes, may be based on user interaction.

[0056] In some implementations, an entity type node may include or be connected to data about: a list of properties associated with that entity type node, the domain to which that entity type node belongs, descriptions, values, any other suitable information, or any combination thereof. A domain refers to a collection of related entity types. For example, the domain [Film] may include, for example, the entity types [Actor], [Director], [Filming Location], [Movie], any other suitable entity type, or any combination thereof. In some implementations, entities are associated with types in more than one domain. For example, the entity node [Benjamin Franklin] may be connected with the entity type node [Politician] in the domain [Government] as well as the entity type node [Inventor] in the domain [Business].

[0057] In some implementations, a node may include or connect to data defining one or more attributes. These may be referred to as attribute references and/or properties. The attribute references may define a particular characteristic of the node. The particular attribute references of a node may depend on what the node represents. In some implementations, an entity reference node may include or connect to: attribute references describing the entity reference, a unique identification reference, a list of entity types associated with the node, a list of differentiation aliases for the node, data associated with the entity reference, a textual description of the entity reference, links to a textual description of the entity reference, other suitable information, or any combination thereof. As described above, nodes may contain a reference or link to long text strings and other information stored in one or more documents external to the data graph. In some implementations, the storage technique may depend on the particular information. For example, a unique identification reference may be stored within the node, a short information string may be stored in a terminal node as a literal, and a

long description of an entity may be stored in an external document linked to by a reference in the data graph.

[0058] Specific values, in some implementations referred to as literals, may be associated with a particular entity in a terminal node by an edge defining the relationship. Literals may refer to values and/or strings of information. For example, literals may include dates, names, and/or numbers. In an example, the entity node [San Francisco] may be connected to a terminal node containing the literal [813000] by an edge annotated with the property [Has Population]. In some implementations, terminal nodes may contain a reference or link to long text strings and other information stored in one or more documents external to the knowledge graph. In some implementations, literals are stored as nodes in the knowledge graph. In some implementations, literals are stored in the knowledge graph but are not assigned a unique identification reference as described below, and are not capable of being associated with multiple entities. In some implementations, literal type nodes may define a type of literal, for example [Date/Time], [Number], or [GPS Coordinates].

[0059] In some implementations, nodes and edges define the relationship between an entity type node and its properties, thus defining a schema. For example, an edge may connect an entity type node to a node associated with a property, which may be referred to as a property node. Entities of the type may be connected to nodes defining particular values of those properties. For example, the entity type node [Person] may be connected to property node [Date of Birth] and a node [Height]. Further, the node [Date of Birth] may be connected to the literal type node [Date/Time], indicating that literals associated with [Date of Birth] include date/time information. The entity node [George Washington], which is connected to entity type node [Person] by an [Is A] edge, may also be connected to a literal [Feb 22, 1732] by the edge [Has Date Of Birth]. In some implementations, the entity node [George Washington] is connected to a [Date Of Birth] property node. It will be understood that in some implementations, both schema and data are modeled and stored in a knowledge graph using the same technique. In this way, both schema and data can be accessed by the same search techniques. In some implementations, schemas are stored in a separate table, graph, list, other data structure, or any combination thereof. It will also be understood that properties may be modeled by nodes, edges, literals, any other suitable data, or any combination thereof.

[0060] For example, the entity node [George Washington] may be connected by an [Is A] edge to the entity type node representing [Person], thus indicating an entity type of the entity, and may also be connected to a literal [Feb 22, 1732] by the edge [Has Date Of Birth],

thus defining a property of the entity. In this way, the knowledge graph defines both entity types and properties associated with a particular entity by connecting to other nodes. In some implementations, [Feb 22, 1732] may be a node, such that it is connected to other events occurring on that date. In some implementations, the date may be further connected to a year node, a month node, and a day of node. It will be understood that this information may be stored in any suitable combination of literals, nodes, terminal nodes, interconnected entities, any other suitable arrangement, or any combination thereof.

[0061] In some implementations, entity types, properties, and other suitable content is created, defined, redefined, altered, or otherwise generated by any suitable technique. For example, content may be generated by manual user input, by automatic responses to user interactions, by importation of data from external sources, by any other suitable technique, or any combination thereof. For example, if a commonly searched for term is not represented in the knowledge graph, one or more nodes representing that node may be added. In another example, a user may manually add information and organizational structures.

[0062] In some implementations, the knowledge graph may include information for differentiation and disambiguation of terms and/or entities. As used herein, differentiation refers to the many-to-one situation where multiple names are associated with a single entity. As used herein, disambiguation refers to the one-to-many situation where the same name is associated with multiple entities. In some implementations, nodes may be assigned a unique identification reference. In some implementations, the unique identification reference may be an alphanumeric string, a name, a number, a binary code, any other suitable identifier, or any combination thereof. The unique identification reference may allow the system to assign unique references to nodes with the same or similar textual identifiers. In some implementations, the unique identifiers and other techniques are used in differentiation, disambiguation, or both. For example, there may be an entity reference node related to the city [Philadelphia], an entity reference node related to the movie [Philadelphia], and an entity reference node related to the cream cheese brand [Philadelphia]. Each of these nodes may have a unique identification reference, stored for example as a number, for disambiguation within the data graph. In some implementations, disambiguation in the data graph is provided by the connections and relationships between multiple nodes. For example, the city [New York] may be disambiguated from the state [New York] because the city is connected to an entity type [City] and the state is connected to an entity type [State]. It will be understood that more complex relationships may also define and disambiguate nodes. For example, a node may be defined by associated entity types, by other entity references

connected to it by particular properties, by its name, by any other suitable information, or any combination thereof. These connections may be useful in disambiguating, for example, the node [Georgia] that is connected to the node [United States] may be understood represent the U.S. State, while the node [Georgia] connected to the nodes [Asia] and [Eastern Europe] may be understood to represent the country in eastern Europe.

[0063] Knowledge graph portion 800 includes information related to the entity [George Washington], represented by [George Washington] node 802. [George Washington] node 802 is connected to [U.S. President] entity type node 804 by [Is A] edge 814 with the semantic content [Is A], such that the 3-tuple defined by nodes 802 and 804 and the edge 814 contains the information "George Washington is a U.S. President." Similarly, "Thomas Jefferson Is A U.S. President" is represented by the tuple of [Thomas Jefferson] node 810, [Is A] edge 820, and [U.S. President] node 804. Knowledge graph portion 800 includes entity type nodes [Person] 824, and [U.S. President] node 804. The person type is defined in part by the connections from [Person] node 824. For example, the type [Person] is defined as having the property [Date Of Birth] by node 830 and edge 832, and is defined as having the property [Gender] by node 834 and edge 836. These relationships define in part a schema associated with the entity type [Person].

[0064] [George Washington] node 802 is shown in knowledge graph portion 800 to be of the entity types [Person] and [U.S. President], and thus is connected to nodes containing values associated with those types. For example, [George Washington] node 802 is connected by [Has Gender] edge 818 to [Male] node 806, thus indicating that "George Washington has gender Male." Further, [Male] node 806 may be connected to the [Gender] node 834 indicating that "Male Is A Type Of Gender." Similarly, [George Washington] node 802 is be connected by [Has Date of Birth] edge 816 to [Feb 22, 1732] node 808, thus indicating that "George Washington Has Date Of Birth Feb 22, 1732." [George Washington] node 802 may also be connected to [1789] node 828 by [Has Assumed Office Date] edge 812.

[0065] Knowledge graph portion 800 also includes [Thomas Jefferson] node 810, connected by [Is A] edge 820 to entity type [U.S. President] node 804 and by [Is A] edge 828 to [Person] entity type node 824. Thus, knowledge graph portion 800 indicates that "Thomas Jefferson" has the entity types "U.S. President" and "Person." In some implementations, [Thomas Jefferson] node 810 is connected to nodes not shown in FIG. 8 referencing his date of birth, gender, and assumed office date.

[0066] It will be understood that knowledge graph portion 800 is merely an example and that it may include nodes and edges not shown. For example, [U.S. President] node 804 may be connected to all of the U.S. Presidents. [U.S. President] node 804 may also be connected to properties related to the entity type such as a duration of term, for example [4 Years], a term limit, for example [2 Terms], a location of office, for example [Washington D.C.], any other suitable data, or any combination thereof. For example, [U.S. President] node 804 is connected to [Assumed Office Date] node 838 by [Has Property] edge 840, defining in part a schema for the type [U.S. President]. Similarly, [Thomas Jefferson] node 810 may be connected to any suitable number of nodes containing further information related to his illustrated entity type nodes [U.S. President], and [Person], and to other entity type nodes not shown such as [Inventor], [Vice President], and [Author]. In a further example, [Person] node 824 may be connected to all entities in the knowledge graph with the type [Person]. In a further example, [1789] node 828 may be connected to all events in the knowledge graph with the property of year [1789]. [1789] node 828 is unique to the year 1789, and disambiguated from, for example, a book entitled [1789], not shown in FIG. 8, by its unique identification reference. In some implementations, [1789] node 828 is connected to the entity type node [Year].

[0067] FIG. 9 shows illustrative knowledge graph portion 900 in accordance with some implementations of the present disclosure. Knowledge graph portion 900 includes [California] node 402, which may also be associated with differentiation aliases such as, for example, [CA], [Calif.], [Golden State], any other suitable differentiation aliases, or any combination thereof. In some implementations, these differentiations are stored in [California] node 902. California is connected by [Is A] edge 904 to the [U.S. State] entity type node 906. [New York] node 910 and [Texas] node 914 are also connected to [U.S. State] node 906 by [Is A] edges 908 and 912, respectively. [California] node 902 is connected by [Has Capital City] edge 920 to [Sacramento] node 922, indicating the information that "California Has Capital City Sacramento." Sacramento node 922 is further connected by [Is A] edge 924 to the [City] entity type node 926. Similarly, [Texas] node 914 is connected by [Has City] edge 920 to [Houston] node 928, which is further connected to the [City] entity type node 926 by [Is A] edge 840. [California] node 902 is connected by [Has Population] edge 916 to node 918 containing the literal value [37,691,912]. In an example, the particular value [37,691,912] may be periodically automatically updated by the knowledge graph based on an external website or other source of data. Knowledge graph portion 900 may include other nodes not shown. For example, [U.S. State] entity type node

906 may be connected to nodes defining properties of that type such as [Population] and [Capital City]. These type—property relationships may be used to define other relationships in knowledge graph portion 900 such as [Has Population] edge 916 connecting entity node [California] 902 with terminal node 918 containing the literal defining the population of California.

[0068] It will be understood that while knowledge graph portion 800 of FIG. 8 and knowledge graph portion 900 of FIG. 9 show portions of a knowledge graph, all pieces of information may be contained within a single graph and that these selections illustrated herein are merely an example. In some implementations, separate knowledge graphs are maintained for different respective domains, for different respective entity types, or according to any other suitable delimiting characteristic. In some implementations, separate knowledge graphs are maintained according to size constraints. In some implementations, a single knowledge graph is maintained for all entities and entity types.

[0069] A knowledge graph, or any other suitable data structure, may be implemented using any suitable software constructs. In an example, a knowledge graph is implemented using object oriented constructs in which each node is an object with associated functions and variables. Edges, in this context, may be objects having associated functions and variables. In some implementations, data contained in a knowledge graph, pointed to by nodes of a knowledge graph, or both, is stored in any suitable one or more data repositories across one or more servers located in one or more geographic locations coupled by any suitable network architecture.

[0070] FIG. 10 shows a flow diagram 1000 of illustrative steps for providing search results in accordance with some implementations of the present disclosure. Illustrative steps to receive a search query associated with an entity reference 1002, provide a set of results for the search query, where the set of results distinguishes between distinct entities 1004, identify one or more attributes of at least one entity of the one or more distinct entities based on the set of results 1006, generate one or more additional search queries based on the search query, the at least one entity, and the one or more attributes 1008, receive an input selecting at least one or more of the additional search queries 1010, and provide an updated set of results based on the selected one or more additional search queries 1012.

[0071] In some implementations, the system may implement steps 1004, 1006, and 1008 based at least in part on the knowledge graph, or any other suitable data graph, as discussed in detail in relation to FIGS. 7–9. For example, the knowledge graph may be used as a source of data for structured data block 104 of FIG. 1.

[0072] In some implementations, the system may be used to provide search results based on medical related search queries. For example, if a search query is provided to the system that includes an entity reference related to a particular health condition, then the system may provide augmented queries that contain that entity reference plus related attributes. This allows a user to perform more comprehensive searches based on aspects of a medical condition that is more relevant to that user.

[0073] The following description and accompanying FIGS. 11 and 12 describe illustrative computer systems that may be used in some implementations of the present disclosure. It will be understood that elements of FIGS. 11 and 12 are merely exemplary and that any suitable elements may be added, removed, duplicated, replaced, or otherwise modified.

[0074] It will be understood that the system may be implemented on any suitable computer or combination of computers. In some implementations, the system is implemented in a distributed computer system including two or more computers. In an example, the system may use a cluster of computers located in one or more locations to perform processing and storage associated with the system. It will be understood that distributed computing may include any suitable parallel computing, distributed computing, network hardware, network software, centralized control, decentralized control, any other suitable implementations, or any combination thereof.

[0075] FIG. 11 shows illustrative computer system 1100 in accordance with some implementations of the present disclosure. System 1100 may include one or more computing device 1102. In some implementations, computing device 1102, and any other device of system 1100, includes one or more computers and/or one or more processors. In some implementations, a processor includes one or more hardware processors, for example, integrated circuits, one or more software modules, computer-readable media such as memory, firmware, or any combination thereof. In some implementations, computing device 1102 includes one or more computer-readable medium storing software, include instructions for execution by the one or more processors for performing the techniques discussed above with respect to FIGS. 7 and 8, or any other techniques disclosed herein. In some implementations, computing device 1102 includes a smartphone, tablet computer, desktop computer, laptop computer, server, personal digital assistant (PDA), portable audio player, portable video player, mobile gaming device, other suitable user device capable of providing content, or any combination thereof.

[0076] Computing device 1102 may be coupled to network 1104 directly through connection 1006, through wireless repeater 1110, by any other suitable way of coupling to network 1104, or by any combination thereof. Network 1104 may include the Internet, a dispersed network of computers and servers, a local network, a public intranet, a private intranet, other coupled computing systems, or any combination thereof.

[0077] Computing device 1102 may be coupled to network 1104 by wired connection 1106. Connection 1106 may include Ethernet hardware, coaxial cable hardware, DSL hardware, T-1 hardware, fiber optic hardware, analog phone line hardware, any other suitable wired hardware capable of communicating, or any combination thereof. Connection 1106 may include transmission techniques including TCP/IP transmission techniques, IEEE 1112 transmission techniques, Ethernet transmission techniques, DSL transmission techniques, fiber optic transmission techniques, ITU-T transmission techniques, any other suitable transmission techniques, or any combination thereof.

[0078] Computing device 1102 may be wirelessly coupled to network 1104 by wireless connection 1108. In some implementations, wireless repeater 1110 receives transmitted information from computing device 1102 by wireless connection 1108 and communicates it with network 1104 by connection 1112. Wireless repeater 1110 receives information from network 1104 by connection 1012 and communicates it with computing device 1102 by wireless connection 1108. In some implementations, wireless connection 1108 may include cellular phone transmission techniques, code division multiple access (CDMA) transmission techniques, global system for mobile communications (GSM) transmission techniques, general packet radio service (GPRS) transmission techniques, , satellite transmission techniques, infrared transmission techniques, Bluetooth transmission techniques, Wi-Fi transmission techniques, WiMax transmission techniques, any other suitable transmission techniques, or any combination thereof.

[0079] Connection 1112 may include Ethernet hardware, coaxial cable hardware, DSL hardware, T-1 hardware, fiber optic hardware, analog phone line hardware, wireless hardware, any other suitable hardware capable of communicating, or any combination thereof. Connection 1012 may include wired transmission techniques including TCP/IP transmission techniques, IEEE 1112 transmission techniques, Ethernet transmission techniques, DSL transmission techniques, fiber optic transmission techniques, ITU-T transmission techniques, any other suitable transmission techniques, or any combination thereof. Connection 1012 may include may include wireless transmission techniques including cellular phone transmission techniques, code division multiple access (CDMA)

transmission techniques, global system for mobile communications (GSM) transmission techniques, general packet radio service (GPRS) transmission techniques, satellite transmission techniques, infrared transmission techniques, Bluetooth transmission techniques, Wi-Fi transmission techniques, WiMax transmission techniques, any other suitable transmission techniques, or any combination thereof.

[0080] Wireless repeater 1110 may include any number of cellular phone transceivers, network routers, network switches, communication satellites, other devices for communicating information from computing device 1102 to network 1104, or any combination thereof. It will be understood that the arrangement of connection 1106, wireless connection 1108 and connection 1112 is merely illustrative and that system 1100 may include any suitable number of any suitable devices coupling computing device 1102 to network 1104. It will also be understood that any computing device 1102, may be communicatively coupled with any user device, remote server, local server, any other suitable processing equipment, or any combination thereof, and may be coupled using any suitable technique as described above.

[0081] In some implementations, any suitable number of remote servers 1114, 1116, 1118, 1120, may be coupled to network 1104. Remote servers may be general purpose, specific, or any combination thereof. One or more search engine servers 1122 may be coupled to the network 1104. In some implementations, search engine server 1122 may include the knowledge graph, may include processing equipment configured to access the knowledge graph, may include processing equipment configured to receive search queries related to the knowledge graph, may include any other suitable information or equipment, or any combination thereof. One or more database servers 1124 may be coupled to network 1104. In some implementations, database server 1124 may store the knowledge graph. In some implementations, where there is more than one knowledge graph, the more than one may be included in database server 1124, may be distributed across any suitable number of database servers and general purpose servers by any suitable technique, or any combination thereof. It will also be understood that the system may use any suitable number of general purpose, specific purpose, storage, processing, search, any other suitable server, or any combination.

[0082] FIG. 12 includes a block diagram of illustrative computing device 1200 in accordance with some implementations of the present disclosure. FIG. 11 includes computing device 1200. In some implementations, computing device 1200 corresponds to computing device 1102 of FIG. 11, a remote computer illustrated in system 1100 of FIG. 11,

any other suitable computer corresponding to system 1100 of FIG. 11, any other suitable device, or any combination thereof. In some implementations, computing device 1210 is an illustrative local and/or remote computer that is part of a distributed computing system. Computing device 1210 may include input/output equipment 1212 and processing equipment 1204. Input/output equipment 1212 may include display 1206, touchscreen 1208, button 1210, accelerometer 1212, global positions system (GPS) receiver 1236, camera 1238, keyboard 1040, mouse 1042, and audio equipment 1234 including speaker 1214 and microphone 1216. In some implementations, the equipment of computing device 1200 may be representative of equipment included in a smartphone user device. It will be understood that the specific equipment included in the illustrative computer system may depend on the type of user device. For example, the input/output equipment 1212 of a desktop computer may include a keyboard 1040 and mouse 1042 and may omit accelerometer 1212 and GPS receiver 1236. It will be understood that computing device 1200 may omit any suitable illustrated elements, and may include equipment not shown such as media drives, data storage, communication devices, display devices, processing equipment, any other suitable equipment, or any combination thereof.

[0083] In some implementations, display 1206 may include a liquid crystal display, light emitting diode display, organic light emitting diode display, amorphous organic light emitting diode display, plasma display, cathode ray tube display, projector display, any other suitable type of display capable of displaying content, or any combination thereof. Display 1206 may be controlled by display controller 1218 or by processor 1224 in processing equipment 1204, by processing equipment internal to display 1206, by other controlling equipment, or by any combination thereof. In some implementations, display 1206 may display data from a knowledge graph.

[0084] Touchscreen 1208 may include a sensor capable of sensing pressure input, capacitance input, resistance input, piezoelectric input, optical input, acoustic input, any other suitable input, or any combination thereof. Touchscreen 1208 may be capable of receiving touch-based gestures. Received gestures may include information relating to one or more locations on the surface of touchscreen 1208, pressure of the gesture, speed of the gesture, duration of the gesture, direction of paths traced on its surface by the gesture, motion of the device in relation to the gesture, other suitable information regarding a gesture, or any combination thereof. In some implementations, touchscreen 1208 may be optically transparent and located above or below display 1206. Touchscreen 1208 may be coupled to and controlled by display controller 1218, sensor controller 1220, processor 1224, any other

suitable controller, or any combination thereof. In some implementations, touchscreen 1208 may include a virtual keyboard capable of receiving, for example, a search query used to identify data in a knowledge graph.

[0085] In some embodiments, a gesture received by touchscreen 1208 may cause a corresponding display element to be displayed substantially concurrently (i.e., immediately following or with a short delay) by display 1206. For example, when the gesture is a movement of a finger or stylus along the surface of touchscreen 1208, the system may cause a visible line of any suitable thickness, color, or pattern indicating the path of the gesture to be displayed on display 1206. In some implementations, for example, a desktop computer using a mouse, the functions of the touchscreen may be fully or partially replaced using a mouse pointer displayed on the display screen.

[0086] Button 1210 may be one or more electromechanical push-button mechanism, slide mechanism, switch mechanism, rocker mechanism, toggle mechanism, other suitable mechanism, or any combination thereof. Button 1210 may be included in touchscreen 1208 as a predefined region of the touchscreen (e.g., soft keys). Button 1210 may be included in touchscreen 1208 as a region of the touchscreen defined by the system and indicated by display 1206. Activation of button 1210 may send a signal to sensor controller 1220, processor 1224, display controller 1220, any other suitable processing equipment, or any combination thereof. Activation of button 1210 may include receiving from the user a pushing gesture, sliding gesture, touching gesture, pressing gesture, time-based gesture (e.g., based on the duration of a push), any other suitable gesture, or any combination thereof.

[0087] Accelerometer 1212 may be capable of receiving information about the motion characteristics, acceleration characteristics, orientation characteristics, inclination characteristics and other suitable characteristics, or any combination thereof, of computing device 1200. Accelerometer 1212 may be a mechanical device, microelectromechanical (MEMS) device, nanoelectromechanical (NEMS) device, solid state device, any other suitable sensing device, or any combination thereof. In some implementations, accelerometer 1212 may be a 3-axis piezoelectric microelectromechanical integrated circuit which is configured to sense acceleration, orientation, or other suitable characteristics by sensing a change in the capacitance of an internal structure. Accelerometer 1212 may be coupled to touchscreen 1208 such that information received by accelerometer 1212 with respect to a gesture is used at least in part by processing equipment 1204 to interpret the gesture.

[0088] Global positioning system (GPS) receiver 1236 may be capable of receiving signals from global positioning satellites. In some implementations, GPS receiver 1236 may

receive information from one or more satellites orbiting the earth, the information including time, orbit, and other information related to the satellite. This information may be used to calculate the location of computing device 1200 on the surface of the earth. GPS receiver 1236 may include a barometer (not shown) to improve the accuracy of the location. GPS receiver 1236 may receive information from other wired and wireless communication sources regarding the location of computing device 1200. For example, the identity and location of nearby cellular phone towers may be used in place of, or in addition to, GPS data to determine the location of computing device 1200.

[0089] Camera 1238 may include one or more sensors to detect light. In some implementations, camera 1238 may receive video images, still images, or both. Camera 1238 may include a charged coupled device (CCD) sensor, a complementary metal oxide semiconductor (CMOS) sensor, a photocell sensor, an IR sensor, any other suitable sensor, or any combination thereof. In some implementations, camera 1238 may include a device capable of generating light to illuminate a subject, for example, an LED light. Camera 1238 may communicate information captured by the one or more sensor to sensor controller 1220, to processor 1224, to any other suitable equipment, or any combination thereof. Camera 1238 may include lenses, filters, and other suitable optical equipment. It will be understood that computing device 1200 may include any suitable number of camera 1238.

[0090] Audio equipment 1234 may include sensors and processing equipment for receiving and transmitting information using acoustic or pressure waves. Speaker 1214 may include equipment to produce acoustic waves in response to a signal. In some implementations, speaker 1214 may include an electroacoustic transducer wherein an electromagnet is coupled to a diaphragm to produce acoustic waves in response to an electrical signal. Microphone 1216 may include electroacoustic equipment to convert acoustic signals into electrical signals. In some implementations, a condenser-type microphone may use a diaphragm as a portion of a capacitor such that acoustic waves induce a capacitance change in the device, which may be used as an input signal by computing device 1200.

[0091] Speaker 1214 and microphone 1216 may be contained within computing device 1200, may be remote devices coupled to computing device 1200 by any suitable wired or wireless connection, or any combination thereof.

[0092] Speaker 1214 and microphone 1216 of audio equipment 1234 may be coupled to audio controller 1222 in processing equipment 1204. This controller may send and receive signals from audio equipment 1234 and perform pre-processing and filtering steps before

transmitting signals related to the input signals to processor 1224. Speaker 1214 and microphone 1216 may be coupled directly to processor 1224. Connections from audio equipment 1234 to processing equipment 1204 may be wired, wireless, other suitable arrangements for communicating information, or any combination thereof.

[0093] Processing equipment 1204 of computing device 1200 may include display controller 1218, sensor controller 1220, audio controller 1222, processor 1224, memory 1226, communication controller 1228, and power supply 1232.

[0094] Processor 1224 may include circuitry to interpret signals input to computing device 1200 from, for example, touchscreen 1208 and microphone 1216. Processor 1224 may include circuitry to control the output to display 1206 and speaker 1214. Processor 1224 may include circuitry to carry out instructions of a computer program. In some implementations, processor 1224 may be an integrated electronic circuit based, capable of carrying out the instructions of a computer program and include a plurality of inputs and outputs.

[0095] Processor 1224 may be coupled to memory 1226. Memory 1226 may include random access memory (RAM), flash memory, programmable read only memory (PROM), erasable programmable read only memory (EPROM), magnetic hard disk drives, magnetic tape cassettes, magnetic floppy disks optical CD-ROM discs, CD-R discs, CD-RW discs, DVD discs, DVD+R discs, DVD-R discs, any other suitable storage medium, or any combination thereof.

[0096] The functions of display controller 1218, sensor controller 1220, and audio controller 1222, as have been described above, may be fully or partially implemented as discrete components in computing device 1200, fully or partially integrated into processor 1224, combined in part or in full into combined control units, or any combination thereof.

[0097] Communication controller 1228 may be coupled to processor 1224 of computing device 1200. In some implementations, communication controller 1228 may communicate radio frequency signals using antenna 1230. In some implementations, communication controller 1228 may communicate signals using a wired connection, not shown. Wired and wireless communications communicated by communication controller 1228 may use Ethernet, amplitude modulation, frequency modulation, bitstream, code division multiple access (CDMA), global system for mobile communications (GSM), general packet radio service (GPRS), satellite, infrared, Bluetooth, Wi-Fi, WiMax, any other suitable communication configuration, or any combination thereof. The functions of communication controller 1228 may be fully or partially implemented as a discrete component in computing

device 1200, may be fully or partially included in processor 1224, or any combination thereof. In some implementations, communication controller 1228 may communicate with a network such as network 1004 of FIG. 10 and may receive information from a knowledge graph stored, for example, in database 1024 of FIG. 10.

[0098] Power supply 1232 may be coupled to processor 1224 and to other components of computing device 1200. Power supply 1232 may include a lithium-polymer battery, lithium-ion battery, NiMH battery, alkaline battery, lead-acid battery, fuel cell, solar panel, thermoelectric generator, any other suitable power source, or any combination thereof. Power supply 1232 may include a hard wired connection to an electrical power source, and may include electrical equipment to convert the voltage, frequency, and phase of the electrical power source input to suitable power for computing device 1200. In some implementations of power supply 1232, a wall outlet may provide 1020 volts, 60Hz AC. A circuit of transformers, resistors, inductors, capacitors, transistors, and other suitable electronic components included in power supply 1232 may convert the 1020V AC from a wall outlet power to 5 volts DC. In some implementations of power supply 1232, a lithium-ion battery including a lithium metal oxide-based cathode and graphite-based anode may supply 3.7V to the components of computing device 1200. Power supply 1232 may be fully or partially integrated into computing device 1200, or may function as a stand-alone device. Power supply 1232 may power computing device 1200 directly, may power computing device 1200 by charging a battery, may provide power by any other suitable way, or any combination thereof.

[0099] The foregoing is merely illustrative of the principles of this disclosure and various modifications may be made by those skilled in the art without departing from the scope of this disclosure. The above described implementations are presented for purposes of illustration and not of limitation. The present disclosure also may take many forms other than those explicitly described herein. Accordingly, it is emphasized that this disclosure is not limited to the explicitly disclosed methods, systems, and apparatuses, but is intended to include variations to and modifications thereof, which are within the spirit of the following claims.

What is claimed is:

1. A computer-implemented method comprising:
 - receiving a search query associated with an entity reference, wherein the entity reference corresponds to one or more distinct entities;
 - providing a set of results for the search query wherein the set of results distinguishes between the one or more distinct entities;
 - identifying one or more attributes of at least one entity of the one or more distinct entities based at least in part on the set of results;
 - generating one or more additional search queries based on the search query, the at least one entity, and the one or more attributes;
 - receiving an input selecting at least one of the one or more additional search queries; and
 - providing an updated set of results based on the selected one or more additional search queries, wherein the updated set of results comprises at least one result not in the set of results.
2. The method of claim 1, wherein generating one or more additional search queries comprises:
 - ranking the identified one or more attributes; and
 - generating one or more additional search queries based on the search query, the at least one entity, the one or more attributes, and the ranking.
3. The method of claim 2, wherein the ranking is based on frequency of occurrence.
4. The method of claim 2, wherein the ranking is based on a location of each of the one or more attributes with respect to the entity reference in the set of results.
5. The method of claim 1, further comprising displaying the one or more additional search queries.
6. The method of claim 1, wherein identifying one or more attributes comprises identifying one or more attributes from a structured data source.

7. The method of claim 6, wherein the structured data source is a social graph.
8. The method of claim 6, wherein the structured data source is a knowledge graph.
9. The method of claim 1, wherein identifying one or more attributes comprises identifying one or more attributes from unstructured data.
10. A system comprising one or more computers configured to perform operations comprising:
 - receiving a search query associated with an entity reference, wherein the entity reference corresponds to one or more distinct entities;
 - providing a set of results for the search query wherein the set of results distinguishes between distinct entities;
 - identifying one or more attributes of at least one entity of the one or more distinct entities based at least in part on the set of results
 - generating one or more additional search queries based on the search query, the at least one entity, and the one or more attributes;
 - receiving an input selecting at least one of the one or more additional search queries; and
 - providing an updated set of results based on the selected one or more additional search queries, wherein the updated set of results comprises at least one result not in the set of results.
11. The system of claim 10, wherein the one or more computers are further configured to perform operations comprising:
 - ranking the identified one or more attributes; and
 - generating one or more additional search queries based on the search query, the at least one entity, the one or more attributes, and the ranking.
12. The system of claim 10, wherein the ranking is based on frequency of occurrence.
13. The system of claim 10, wherein the ranking is based on a location of each of the one or more attributes with respect to the entity reference in the set of results.

14. The system of claim 7, wherein the one or more computers are further configured to perform operations comprising causing to be displayed on a user device the one or more additional search queries.
15. The system of claim 10, wherein the one or more computers are further configured to perform operations comprising identifying one or more attributes from a structured data source.
16. The system of claim 15, wherein the structured data source is a social graph.
17. The system of claim 15, wherein the structured data source is a knowledge graph.
18. The system of claim 10, wherein the one or more computers are further configured to perform operations comprising identifying one or more attributes from unstructured data.
19. A computer-readable medium storing instructions that, when executed by one or more processors, cause the one or more processors to perform operations comprising:
 - receiving a search query associated with an entity reference, wherein the entity reference corresponds to one or more distinct entities;
 - providing a set of results for the search query wherein the set of results distinguishes between the one or more distinct entities;
 - identifying one or more attributes of at least one entity of the one or more distinct entities based at least in part on the set of results;
 - generating one or more additional search queries based on the search query, the at least one entity, and the one or more attributes;
 - receiving an input selecting at least one of the one or more additional search queries; and
 - providing an updated set of results based on the selected one or more additional search queries, wherein the updated set of results comprises at least one result not in the set of results.
20. The computer-readable medium of claim 19, wherein generating one or more additional search queries comprises:
 - ranking the identified one or more attributes; and

generating one or more additional search queries based on the search query, the at least one entity, the one or more attributes, and the ranking.

21. The computer-readable medium of claim 20, wherein the ranking is based on frequency of occurrence.
22. The computer-readable medium of claim 20, wherein the ranking is based on a location of each of the one or more attributes with respect to the entity reference in the set of results.
23. The computer-readable medium of claim 19, further comprising displaying the one or more additional search queries.
24. The computer-readable medium of claim 19, wherein identifying one or more attributes comprises identifying one or more attributes from a structured data source.
25. The computer-readable medium of claim 24, wherein the structured data source is a social graph.
26. The computer-readable medium of claim 24, wherein the structured data source is a knowledge graph.
27. The computer-readable medium of claim 19, wherein identifying one or more attributes comprises identifying one or more attributes from unstructured data.

1/12

100




FIG. 1

2/12

200




FIG. 2

3/12


300

FIG. 3

4/12




FIG. 4

5/12




FIG. 5

6/12




FIG. 6

7/12

700




FIG. 7

8/12




FIG. 8

9/12




FIG. 9

10/12




FIG. 10

11/12


1100

FIG. 11

12/12




FIG. 12

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/024271

A. CLASSIFICATION OF SUBJECT MATTER

INV. G06F17/30

ADD.

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)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2009/171929 A1 (JING FENG [CN] ET AL) 2 July 2009 (2009-07-02) abstract; figures 2-3,7-9 paragraph [0040] - paragraph [0046] paragraph [0050] - paragraph [0059] paragraph [0092] - paragraph [0094] -----	1-27
X	US 2011/299743 A1 (KE QIFA [US] ET AL) 8 December 2011 (2011-12-08)	1,10,19
Y	abstract paragraph [0032] - paragraph [0035] ----- -/-	2-9, 11-18, 20-27



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

11 June 2014

Date of mailing of the international search report

18/06/2014

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Nazzaro, Antonio

INTERNATIONAL SEARCH REPORT

International application No

PCT/US2014/024271

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	RABIA NURAY-TURAN ET AL: "Exploiting Web querying for Web people search", ACM TRANSACTIONS ON DATABASE SYSTEMS, vol. 37, no. 1, 1 February 2012 (2012-02-01), pages 1-41, XP055122694, ISSN: 0362-5915, DOI: 10.1145/2109196.2109203	2-9, 11-18, 20-27
A	abstract page 7.4 - page 7.9	1,10,19
A	----- Man Lan ET AL: "Which Who are They? People Attribute Extraction and Disambiguation in Web Search Results *", Proceedings WWW 2009, 20 April 2009 (2009-04-20), XP055122693, Madrid, Spain Retrieved from the Internet: URL:https://137.132.80.57/~tancl/publicati ons/c2009/WePS2009FullDescription-Final.pd f [retrieved on 2014-06-11] the whole document -----	1-27

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2014/024271

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009171929	A1	02-07-2009	NONE

US 2011299743	A1	08-12-2011	NONE
