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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
5 July 2012 (05.07.2012)

WIPO | PCT

(10) International Publication Number
WO 2012/088591 A1

(51) International Patent Classification:
H04L 12/16 (2006.01) *G06F 17/30* (2006.01)

(21) International Application Number:
PCT/CA2011/001403

(22) International Filing Date:
30 December 2011 (30.12.2011)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:

61/428,676	30 December 2010 (30.12.2010)	US
61/428,435	30 December 2010 (30.12.2010)	US
61/428,445	30 December 2010 (30.12.2010)	US
61/428,607	30 December 2010 (30.12.2010)	US
61/430,090	5 January 2011 (05.01.2011)	US
61/430,138	5 January 2011 (05.01.2011)	US
61/430,141	5 January 2011 (05.01.2011)	US
61/430,143	5 January 2011 (05.01.2011)	US
13/162,069	16 June 2011 (16.06.2011)	US
13/340,820	30 December 2011 (30.12.2011)	US
13/340,792	30 December 2011 (30.12.2011)	US

(71) Applicant *(for all designated States except US)*: **PRIMAL FUSION INC.** [CA/CA]; 7-258 King Street North, Waterloo, Ontario N2J 2Y9 (CA).

(72) Inventors; and

(75) Inventors/Applicants *(for US only)*: **SWEENEY, Peter** [CA/CA]; 56 Kilbirnie Court, Kitchener, Ontario N2R 1B8 (CA). **ILYAS, Ihab, F.** [CA/CA]; 294 Wiltshire Place, Waterloo, Ontario N2T 1 R4 (CA). **KHAN, Naim** [CA/CA]; 22 Bridge Street East, Kitchener, Ontario N2K 1J5 (CA).

(74) Agent: **NAHM, Tai, W.**; c/o Miller Thomson LLP, Accelerator Building, 295 hagey Boulevard, Suite 300, Waterloo, Ontario N2L 6R5 (CA).

(81) Designated States *(unless otherwise indicated, for every kind of national protection available)*: AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR PERFORMING A SEMANTIC OPERATION ON A DIGITAL SOCIAL NETWORK

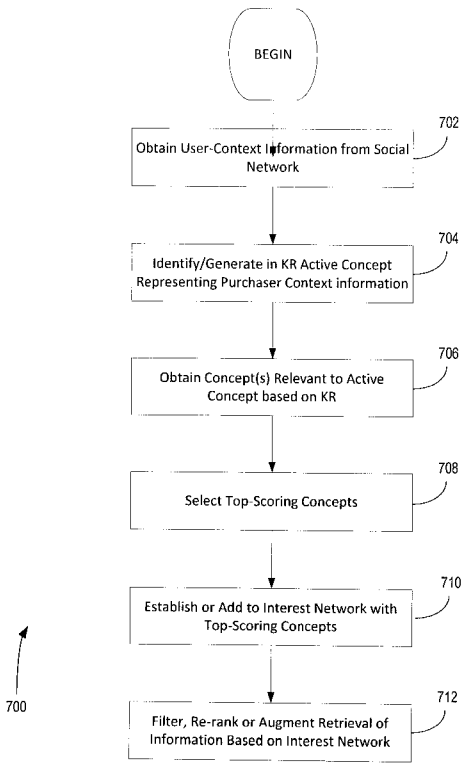

FIG. 7

(57) Abstract: Disclosed is a system and method for performing a semantic operation on a social network. In an embodiment, the method comprises receiving a social network user context associated with a user of the social network; generating, through a semantic operation, an interest network based on the user context information; and filtering, ranking or augmenting, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network; wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network. In another embodiment, the method further comprises representing the interest network as an interest graph. In yet another embodiment, the semantic operation is a synthesis operation or retrieval operation performed on a knowledge representation.

WO 2012/088591 A1

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Published:

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

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

SYSTEM AND METHOD FOR PERFORMING A SEMANTIC OPERATION ON A DIGITAL SOCIAL NETWORK

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] The present application also builds upon concepts disclosed in a number of prior applications by the same inventors and/or assignee, including the following to which the reader is referred for background additional to that discussed below: U.S. Patent Application Ser. No. 13/162,069 filed on June 16, 2011, titled "Methods and Apparatus for Searching of Content Using Semantic Synthesis," attorney docket No. P0913.70013US01; U.S. Patent Application Ser. No. 12/671,846 filed on February 2, 2010, titled "Method System, and Computer Program for User-Driven Dynamic Generation of Semantic Networks and Media Synthesis," attorney docket No. P0913.70007US00; and International Application No. PCT/CA2009/000567 filed May 1, 2009, titled "Method, System, and Computer Program for User-Driven Dynamic Generation of Semantic Networks and Media Synthesis."

FIELD OF INVENTION

[0002] The teachings disclosed herein relate to the field of information retrieval. In particular, the teachings disclosed herein relate to the deployment of systems and methods in a digital information system environment for using information associated with a user in a social network to identify and provide information, from a larger set of digital content, that may be of interest to the user.

BACKGROUND

[0003] Information technology is often used to provide users with various types of information, such as text, audio, video, and any suitable other type of information. In some cases, information is provided to a user in response to an action that the user has taken. For example, information may be provided to a user in response to a search query input by the user or in response to the user having subscribed to content such as an e-mail alert(s) or a electronic newsletter(s). In other cases, information is provided or "pushed" to a user without the user

having specifically requested such information. For example, a user may occasionally be presented with advertisements or solicitations.

[0004] There is a vast array of content that can be provided to users via information technology. Indeed, because of the enormous volume of information available via the Internet, the World Wide Web (WWW), and any other suitable information provisioning sources, and because the available information is distributed across an enormous number of independently owned and operated networks and servers, locating information of interest to users presents challenges. Similar challenges exist when the information of interest is distributed across large private networks.

[0005] Search engines have been developed to aid users in locating desired content on the Internet. A search engine is a computer program that receives a search query from a user (e.g., in the form of a set of keywords) indicative of content desired by the user, and returns information and/or hyperlinks to information that the search engine determines to be relevant to the user's search query.

[0006] Search engines typically work by retrieving a large number of WWW web pages and/or other content using a computer program called a "web crawler" that explores the WWW in an automated fashion (e.g., following every hyperlink that it comes across in each web page that it browses). The located web pages and/or content are analyzed and information about the web pages or content is stored in an index. When a user or an application issues a search query to the search engine, the search engine uses the index to identify the web pages and/or content that it determines to best match the user's search query and returns a list of results with the best-matching web pages and/or content. Frequently, this list is in the form of one or more web pages that include a set of hyperlinks to the web pages and/or content determined to best match the user's search query.

[0007] The sheer volume of content accessible via digital information systems presents a number of information retrieval problems. One challenge is how advertisers can achieve better return on their investment given the vast number of potential users that they could potentially target with advertisements that are relevant to the vast range of interests of such users.

SUMMARY

[0008] The present disclosure relates to a system and method for using information associated with a user in a in a digital information system environment, such as a digital social network, together with one or more data sets expressed as knowledge representations in order to identify and provide information, from a larger set of digital content, that may be of interest to the user.

[0009] In an embodiment, information about a user's online interactions is collected from a number of different sources to create a user context based on the online interactions. The collected information is analyzed to create a comprehensive user context, which may then be used to deliver semantically relevant information to the user.

[0010] In addition to a user's online interaction profile, profiles may also be created for a set or subset of users who are members in on or more digital social networks. The profiles of the users may be overlapped to determine points of intersection between users, whereby relevant information may also be made available to a set or subset of users who are members of an online community.

[0011] Thus, in an aspect, there is provided a method for performing a semantic operation on a social network, the method comprising receiving a social network user context associated with a user of the social network; generating, through a semantic operation, an interest network based on the user context information; and filtering, ranking or augmenting, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network; wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network.

[0012] In an embodiment, the method further comprises representing the interest network as an interest graph.

[0013] In another embodiment, the semantic operation is a synthesis operation or retrieval operation performed on a knowledge representation.

[0014] In another aspect, there is provided a system for performing a semantic operation on a social network, the system adapted to receive a social network user context associated with a user of the social network; generate, through a semantic operation, an interest network based on the user context information; and filter, rank or augment, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network; wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network.

[0015] In yet another aspect, there is provided a non-transitory computer-readable medium storing computer code that when executed on a computer device adapts the device to perform a semantic operation on a social network, the computer-readable medium comprising: code for receiving a social network user context associated with a user of the social network; code for generating, through a semantic operation, an interest network based on the user context information; and code for filtering, ranking or augmenting, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network; wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network.

[0016] In this respect, before explaining at least one embodiment of the system and method of the present disclosure in detail, it is to be understood that the present system and method is not limited in its application to the details of construction and to the arrangements of the components set forth in the following description or illustrated in the drawings. The present system and method is capable of other embodiments and of being practiced and carried out in various ways. Also, it is to be understood that the phraseology and terminology employed herein are for the purpose of description and should not be regarded as limiting.

BRIEF DESCRIPTION OF DRAWINGS

[0017] The accompanying drawings are not intended to be drawn to scale. Like elements are identified by the same or like reference designations when practical. For purposes of clarity, not every component may be labelled in every drawing. In the drawings:

[0018] FIG. 1 is a flowchart of an illustrative process for providing a user with information selected from a large set of digital content, in accordance with some embodiments of the present disclosure.

[0019] FIG. 2 is a block diagram of an illustrative client/server architecture that may be used to implement some embodiments of the present disclosure.

[0020] FIG. 3 is a flowchart of an illustrative process for identifying or generating an active concept representing user context information, in accordance with some embodiments of the present disclosure.

[0021] FIG. 4A-4C is an illustration of generating an active concept representing user context information, in accordance with some embodiments of the present disclosure.

[0022] FIGS. 5A-5H illustrate various approaches for obtaining concepts relevant to an active concept representing user context information, in accordance with some embodiments of the present disclosure.

[0023] FIGS. 6A-6B illustrate techniques for scoring concepts relevant to an active concept, in accordance with some embodiments of the present disclosure.

[0024] FIG. 7 illustrates a process for performing a semantic operation on a digital social network in accordance with an embodiment.

[0025] FIG. 8 illustrates one non-limiting example where the process of FIG. 7may be carried out to establish an interest network for a user.

[0026] FIG. 9 is a block diagram of a computing device on which some embodiments of the present disclosure may be implemented.

[0027] The foregoing is a non-limiting summary of the invention, which is defined by the attached claims.

DETAILED DESCRIPTION

[0028] As noted above, the present disclosure relates to systems and methods for using information associated with a user or users in a in a digital information system environment such as a digital social network in order to identify and provide information, from a larger set of digital content, that may be of interest to the user(s). More particularly, in an embodiment, the present system and method utilizes information associated with a user or users in a digital social network with one or more data sets expressed as knowledge representations in order to located semantically relevant information based on an analysis of a user's digital social networking context.

[0029] The availability of a large volume of information, from various information-provisioning sources such as the Internet, makes it difficult to determine what information may be of interest to users and should be presented to them. For example, it may be difficult to determine what information (e.g., news, advertisements, updates about other people, etc.) should be presented to a user on a website (e.g., a news website, a web portal, a social networking website, etc.) after the user navigates to the website, but without explicitly specifying what information the user is interested in.

[0030] Even in a scenario, such as online search, where the user provides an explicit indication (e.g., a search query) of what information the user may be interested in, such an indication may not be sufficient to accurately identify the content to present to the user from among the large set of content that may be presented to the user. Most conventional search engines simply perform pattern matching between the literal terms in the user's search query and literal terms in the content indexed by the search engine to determine what pieces of content are relevant to the user's query. As such, when the terms in the user's search query do not match the literal terms in the indexed content, the user may not be provided with the information the user is seeking. Also when a user enters a search query containing a term whose meaning is ambiguous, the user may be provided with information entirely unrelated to the meaning of the term that the user intended. As such, the user may be overwhelmed with information irrelevant to the user's interests.

[0031] In these and other settings, semantic processing techniques may be used to identify information, from among a large set of digital content, that the user is likely to be interested in. In particular, applying semantic processing techniques to information associated with a user may help to identify the information in which the user may be interested. As described in greater detail below, information associated with a user may include, but is not limited to, information about the user (e.g., demographic information, geo-spatial information, the user's browsing history, etc.) and/or any information provided by the user (e.g., a search query or queries provided by the user, a user-constructed online profile, etc.).

[0032] Aspects of semantic processing techniques relate to constructing and using knowledge representations to support machine-based storage, knowledge management, and reasoning systems. Conventional methods and systems exist for utilizing knowledge representations (KRs) constructed in accordance with various types of knowledge representation models, which may be realized as concrete data structures, including structured controlled vocabularies such as taxonomies, thesauri, and faceted classifications; formal specifications such as semantic networks and ontologies; and unstructured forms such as documents based in natural language.

[0033] While it is not intended that the claimed invention be limited to specific knowledge representations, a preferred form is the type of formal specification referred to as a semantic network. Semantic networks are explained in many sources, noteworthy among them being U.S. Application Serial No. 12/671,846, titled "Method, System, And Computer Program For User-Driven Dynamic Generation Of Semantic Networks And Media Synthesis" by Peter Sweeney et al., which is hereby incorporated by reference.

[0034] Semantic networks have a broad utility as a type of knowledge representation. As machine-readable data, they can support a number of advanced technologies, such as artificial intelligence, software automation and agents, expert systems, and knowledge management. Additionally, they can be transformed into various forms of media (i.e., other KRs). In other words, the synthesis or creation of semantic networks can support the synthesis of a broad swath of media to extract additional value from the semantic network.

[0035] In some embodiments, a semantic network may be represented as a data structure embodying a directed graph comprising vertices or nodes that represent concepts, and edges that represent semantic relations between the concepts. The data structure embodying a semantic network may be encoded (i.e., instantiated) in a non-transitory, tangible computer-readable storage medium. As such, a semantic network may be said to comprise one or more concepts. Each such concept may be represented by a data structure storing any data associated with one or more nodes in the semantic network representing the concept. An edge in the directed graph may represent any of different types of relationships between the concepts associated with the two nodes that the edge connects. For instance, the relationship represented by an edge in a semantic network may be a "defined-by" relationship or an "is-a" relationship. In the drawings (e.g., FIGS. 5A-5H) that show illustrative semantic networks, "defined-by" relationships are indicated by edges ending with a filled-in circle and "is-a" relationships are indicated by edges ending with an arrow.

[0036] Concepts may be defined in terms of compound levels of abstraction through their relationships to other entities and structurally in terms of other, more fundamental KR entities such as keywords and morphemes. Such a structure may be referred to as a concept definition. Collectively, the more fundamental knowledge representation entities such as keywords and morphemes that comprise concepts are referred to as attributes of the concept.

[0037] As explained in U.S. Patent Application Serial No. 12/671,846, concepts may be defined in terms of compound levels of abstraction through their relationship to other entities and structurally in terms of other, more fundamental knowledge representation entities such as keywords and morphemes. Such a structure may be referred to as a "concept definition."

[0038] Information associated with a user may be used together with at least one knowledge representation, such as a semantic network, in order to infer what information may be of interest to the user. Any suitable knowledge representation may be used. For example, a semantic network representing knowledge associated with content, a subset of which may be of interest to the user, may be used. Such a KR may be constructed from any of numerous data sources including any suitable information provisioning sources and vocabularies.

[0039] Further, information associated with the user may be used together with one or more knowledge representations to identify concepts semantically relevant to information associated with the user. In turn, the identified concepts may be used to determine what information may be of interest to the user in a broad range of applications. For example, finding concepts that are semantically relevant to terms in a user's search query may be useful in determining the semantic meaning of the query and what information the user may be seeking. The query may then be augmented with keywords derived from the identified concepts and improved search results may be returned to the user. As another example, concepts identified as being semantically relevant to information contained in a user's online profile may be useful in determining what information (e.g., advertisements, news, etc.) to present to the user when that user is online.

[0040] Accordingly, in some embodiments, methods for identifying concepts semantically relevant to information associated with the user are disclosed. A concept representing at least a portion of the information associated with a user, termed an "active concept," may be identified or synthesized (i.e., generated) and one or more concepts semantically relevant to the active concept may be obtained. In some instances, concepts semantically relevant to the active concept may be identified in a semantic network; but in other instances, concepts relevant to the active concept may be synthesized by using the active concept and at least one other concept in the semantic network. Concepts semantically relevant to the active concept may be identified and/or synthesized based at least in part on the structure of the semantic network.

[0041] In some embodiments, after concepts semantically relevant to the active concept are obtained, the obtained concepts may be scored. The scores may be calculated in any of numerous ways and, for example, may be calculated based at least in part on the structure of (the data structure of the graph embodying) the semantic network. In turn, the scores may be used to select a subset of the concepts semantically relevant to the active concept. Then, the selected concepts may be used to provide the user with information in which the user may be interested.

[0042] In some embodiments, information associated with one or more users may be used to construct a user-specific knowledge representation corresponding to the user(s). The

user-specific knowledge representation may be constructed from information associated with the user(s) and at least one knowledge representation, such as a semantic network. Accordingly, a user-specific knowledge representation may encode information related to the user(s). Any suitable knowledge representation may be used including, but not limited to, a knowledge representation that represents knowledge associated with content, a subset of which the user(s) may be interested in. The resulting user-specific knowledge representation may be used to identify concepts related to information in which the user(s) may be interested.

[0043] Advantageously, in some embodiments, a user-specific knowledge representation may be stored and used multiple times. This allows semantic processing for a user to take advantage of previously-performed semantic processing for the user and/or one or more other users. For example, if a user-specific knowledge representation was generated based on information associated with one user in a group of users (e.g., an employee of company X), the same user-specific knowledge representation may be used to identify concepts semantically relevant to information associated with another user in the same group (e.g., another employee of company X).

[0044] As previously mentioned, it is often the case that, when users search for information, the terms used in their search queries may not literally match the terms appearing in the content being searched; each side expresses the same or similar concepts in different terms. In such cases, the returned search results, if any, may include fewer quality results than actually are available. Consequently, it is often difficult for users to find all the information they need even when the relevant content exists.

[0045] One example of this situation may occur when a user searches for information within a specialized set of content (e.g., content accessible through a particular website or websites, a particular network, a particular portal, etc.) using terms that appear infrequently in the set of content being searched. For instance, a user may search for information via a medical website within specialized content produced by and intended for medical practitioners and researchers. However, because users of the website may not be medically trained, their search queries may not use many of the terms commonly found in medical articles pertaining to the relevant subject matter. Accordingly, only a few of the terms in user-provided search queries, if

any, may appear in the content accessible through the website and many potentially relevant items may be missed. Another example is when users try to search for information in customer support content. Users may not be technically savvy, but yet need to use specific technical terms in order to find the content that will be helpful to them. Many other examples will be apparent to those skilled in the art.

[0046] Some of the above-mentioned shortcomings of conventional search methods may be addressed by using a secondary or "reference" set of content to improve the quality of search results being returned to users that wish to search a primary, or "target," set of content. The primary set of content may be any suitable set of content and, for example, may be content accessible via a particular website (e.g., an e-commerce website, a website of a business, a website providing access to one or more databases, etc.). A secondary or reference set of content may be any suitable set of content and, for example, may be content in any information repository (e.g., Wikipedia, WordNet, etc.), database, or content-provisioning source. Though it should be recognized that these are only examples and that the target set of content and the reference set of content may be any suitable sets of content, as aspects of the present invention are not limited in this respect.

[0047] By way of illustration, in the above example, a reference domain comprising information about diseases commonly known in the public sphere may help to relate terms in users' search queries in a medical website to terms in content accessible through that website. Indeed, it may be easier to relate a user's search query, such as "Flu Virus," to content accessible through the medical website, which may refer to viral diseases only by using official classifications for the associated virus, if the reference set of content includes information that identifies "Orthomyxoviridae" as a family of influenza viruses and that influenza is commonly known as the "flu." Simply put, the reference set of content may serve as a type of "middle-layer" or a "translation or interpretation medium" to aid in translating terms appearing in search queries to terms appearing in the target set of content being searched.

[0048] Semantic processing techniques may be employed in order to use content in a reference set of content to improve the quality of search results being returned to users that wish to search a target set of content. Accordingly, in some embodiments, a reference knowledge

representation as well as a target knowledge representation may be employed. The reference (target) knowledge representation may be any suitable type of knowledge representation, such as a semantic network, and may represent knowledge associated with content in a reference (target) set of content. The reference (target) knowledge representation may be constructed in any suitable way and, for example, may be constructed based at least in part on content in the reference (target) set of content.

[0049] In some embodiments, the reference and target knowledge representations may be merged to create a merged knowledge representation and one or more terms associated with a user's search query (and, optionally, terms appearing in the user context information associated with the user) may be used to identify or synthesize one or more concepts in the merged knowledge representation that are semantically relevant to the term(s) associated with the user's search query and, optionally, terms appearing in the user context information associated with the user. In turn, the obtained concepts may be used to augment the user's search query with terms associated with the obtained concepts prior to conducting the search. Accordingly, concepts obtained based at least in part on the reference knowledge representation may be used to improve the quality of search results returned in response to users' search queries for content in a target set of content.

[0050] It should be appreciated that the various aspects of the present invention described herein may be implemented in any of numerous ways, and are not limited to any particular implementation technique. Examples of specific implementations are described below for illustrative purposes only, but aspects of the invention described herein are not limited to these illustrative implementations. Additionally, unless it clearly appears otherwise from the particular context, it is intended that the various features and steps described herein may be combined in ways other than the specific example embodiments depicted and that such other combinations are within the scope of the disclosure and are contemplated as inventive.

[0051] FIG. 1 is a flow chart of an illustrative process 100 for providing a user with digital content selected from a large set of digital content based on a knowledge representation that may be used in some embodiments. The process of FIG. 1 begins at act 102, where user context information associated with one or more users may be obtained. As described in greater

detail below, user context information may be any suitable information associated with the user(s) and/or provided by the user(s).The manner in which user context information is obtained is also described in greater detail below.

[0052] Process 100 then continues to act 104, where an active concept representing at least a portion of the user context information may be identified in a knowledge representation. The knowledge representation may be any suitable knowledge representation and, in some embodiments, may be a user-specific knowledge representation associated with the user(s). Though, it should be recognized that the knowledge representation is not limited to being a user-specific knowledge representation and may be any other knowledge representation available to process 100. In some embodiments, as part of act 104, an active concept representing the user context information may be generated, for subsequent use in constructing a user-specific knowledge representation comprising the active concept.

[0053] Process 100 then continues to act 106, where one or more concepts semantically relevant to the active concept may be obtained by using the knowledge representation comprising the active concept. (Example relevance measures are discussed below.) Process 100 then continues to act 108, where one or more of the obtained concept(s) may be selected. The concept(s) may be selected based at least in part on a score that one or more of the concept(s) may be assigned by using a relevance measure. Process 100 then proceeds to act 110, where content may be provided to the one or more users based at least in part on the active concept, identified or generated in act 104, and the concept(s) selected in act 108. Such content may be selected from a large set of content by using the active concept and the concept(s) selected in act 108. Each of the acts of the process of FIG. 1 may be performed in any of a variety of ways, and some examples of the ways in which these acts may be performed in various embodiments are described in greater detail below.

[0054] Process 100 and any of its variants may be implemented using hardware, software or any suitable combination thereof. When implemented in software, the software may execute on any suitable processor or collection of processors, whether stand-alone or networked. The software may be stored as processor-executable instructions and configuration parameters; such

software may be stored on one or more non-transitory, tangible computer-readable storage media.

[0055] Software implementing process 100 may be organized in any suitable manner. For example, it may be organized as a software system comprising one or more software modules such that each software module may perform at least a part of one or more acts of process 100. Though, in some embodiments, one or more software modules of such a software system may perform functions not related to acts of process 100, as aspects of the present invention are not limited in this respect.

I. Obtaining User Context Information

[0056] As discussed above, at act 102 of process 100, user context information associated with one or more users may be obtained. User context information may comprise any information that may be used to identify what information the user(s) may be seeking and/or may be interested in. User context information may also comprise information that may be used to develop a model of the user(s) that may be subsequently used to provide those user(s) with information. As such, user context information may include, but is not limited to, any suitable information related to the user(s) that may be collected from any available sources and/or any suitable information directly provided by the user(s).

[0057] In some embodiments, information related to a user may be any suitable information about the user. For example, information related to a user may comprise demographic information (e.g., gender, age group, education level, etc.) associated with the user. As another example, information related to a user may comprise details of the user's Internet browsing history. Such information may comprise a list of one or more websites that the user may have browsed, the time of any such browsing, and/or the place (i.e., geographic location) from where any such browsing occurred. The user's browsing history may further comprise information that the user searched for and any associated browsing information including, but not limited to, the search results the user obtained in response to any such searches.

[0058] As another example, information related to a user may comprise any information that the user has provided via any user interface on the user's computing device or on one or

more websites that the user may have browsed. For instance, information related to a user may comprise any information associated with the user on any website such as a social networking website, job posting website, a blog, a discussion thread, etc. Such information may include, but is not limited to, the user's profile on the website, any information associated with multimedia (e.g., images, videos, etc.) corresponding to the user's profile, and any other information entered by the user on the website. As yet another example, information related to a user may comprise consumer interaction information as described in U.S. Patent Application Serial No. 12/555,293, filed 09/08/2009, and entitled "Synthesizing Messaging Using Content Provided by Consumers," which is incorporated herein by reference.

[0059] In some embodiments, information related to a user may comprise geo-spatial information. For instance, the geo-spatial information may comprise the current location of the user and/or a computing device of the user (e.g., user's home, library in user's hometown, user's work place, a place to which the user has traveled, and/or the geographical location of the user's device as determined by the user's Internet IP address, etc.). Geo-spatial information may include an association between information about the location of the user's computing device and any content that the user was searching or viewing when the user's computing device was at or near that location. In some embodiments, information related to a user may comprise temporal information. For example, the temporal information may comprise the time during which a user was querying or viewing specific content on a computing device. The time may be specified at any suitable scale such as on the scale of years, seasons, months, weeks, days, hours, minutes, seconds, etc.

[0060] Additionally or alternatively, user context information associated with one or more users may comprise information provided by the user(s). Such information may be any suitable information indicative of what information the user(s) may be interested in. For example, user context information may comprise one or more user search queries input by a user into a search engine (e.g., an Internet search engine, a search engine adapted for searching a particular domain such as a corporate intranet, etc.). As another example, user context information may comprise one or more user-specified indicators of the type of information the user may be interested in. A user may provide the indicator(s) in any of numerous ways. The user may type in or speak an indication of his preferences, select one or more options provided

by a website or an application (e.g., select an item from a dropdown menu, check a box, etc.), highlight or otherwise select a portion of the content of interest to the user on a website or in an application, and/or in any other suitable manner. For example, the user may select one or more options on a website to indicate that he wishes to receive news updates related to a certain topic or topics, advertisements relating to one or more types of product(s), information about updates on any of numerous types of websites, newsletters, e-mail digests, etc.

[0061] The user context information may be obtained, in act 102, in any of a variety of possible ways. For example, in some embodiments, the user context information may be provided from a user's client computer to one or more server computers that execute software code that performs process 100. That is, for example, as shown in FIG. 2, a user may operate a client computer 202 that executes an application program 204. Application program 204 may send user context information 206 (e.g., a search query entered by the user into application program 204) to server computer 208, which may be a computer that performs process 100. Thus, server 208 may receive user context information 206 from application program 204 executing on client 202.

[0062] Application program 204 may be any of a variety of types of application programs that are capable of, directly or indirectly, sending information to and receiving information from server 208. For example, in some embodiments, application program 204 may be an Internet or WWW browser, an instant messaging client, or any other suitable application.

[0063] In the example of FIG. 2, application program 204 is shown as sending the user context information directly to server 208. It should be recognized that this is a simplified representation of how the user context information may be sent from client 202 to server 208, and that the user context information need not be sent directly from client 202 to server 208. For example, in some embodiments, the user's search query may be sent to server 208 via a network. The network may be any suitable type of network such as a LAN, WAN, the Internet, or a combination of networks.

[0064] It should also be recognized that receiving user context information from a user's client computer is not a limiting aspect of the present invention as user context information may be obtained in any other suitable way as part of act 102 of process 100. For example, user

context information may be obtained, actively by requesting and/or passively by receiving, from any source with, or with access to, user context information associated with one or more users.

II. Identifying or Generating Active Concept Representing User Context Information

[0065] As discussed above, at act 104 of process 100, an active concept, representing at least a portion of the user context information obtained during act 102, may be either identified in a knowledge representation, which may be a user-specific knowledge representation or any other suitable knowledge representation, or generated and used to construct a user-specific knowledge representation comprising the active concept. Any of a variety of possible techniques may be used to identify or generate an active concept representing user context information. An example of one such technique that may be used in some embodiments is illustrated in process 300 of FIG. 3.

[0066] Process 300 begins at act 301, where a relevant portion of the user context information may be identified. As previously discussed, user context information may comprise any of numerous types of information including, but not limited to, information about a user (e.g., profile of the user on a website, the user's browsing history, etc.) and information provided by a user (e.g., a search query). Accordingly, various portions of the user context information may be used in different scenarios. For example, when a user is searching for information, a relevant portion of the user context information may comprise a user's search query, but also may comprise additional information that may be helpful in searching for the information that the user seeks (e.g., the user's current location, the user's browsing history, etc.). As another example, when presenting a user with one or more advertisements, a relevant portion of the user context information may comprise information indicative of one or more products that the user may have interest in. As another example, when providing a user with news articles (or any other suitable type of content), a relevant portion of the user context information may comprise information indicative of the user's interests. The relevant portion of the user context information obtained (e.g., in act 102) may be identified in any suitable way as the manner in which the relevant portion of the user context information is identified is not a limitation of aspects of the present invention. It should be also recognized that, in some instances, the relevant portion of the user context information may comprise a subset of the user context information, but, in other

embodiments, the relevant portion may comprise all of the user context information, as aspects of the present invention are not limited in this respect.

[0067] The relevant portion of the user context information, identified in act 301, may be represented in any of numerous ways. For example, in some embodiments, the relevant portion of user context information may be represented via one or more alphanumeric strings. An alphanumeric string may comprise any suitable number of characters (including spaces), words, numbers, and/or any of numerous other symbols. An alphanumeric string may, for example, represent a user search query and/or any suitable information indicative of what information the user may be interested in. Though, it should be recognized that any of numerous other data structures may be used to represent user context information and/or any portion thereof.

[0068] Next, process 300 proceeds to decision block 302, where it is determined whether the relevant portion of the user context information associated with a particular user matches a concept in a knowledge representation. Any suitable knowledge representation may be used. In some instances, a user-specific knowledge representation associated with the user or a group of users that includes the user may be used. However, any other suitable knowledge representation may be used and, for example, a knowledge representation not adapted to any particular user or users may be employed.

[0069] In some embodiments, the knowledge representation used may be a semantic network. Though, in other embodiments, any of other numerous types of knowledge representations may be employed (e.g., a non-graphical knowledge representation). The knowledge representation may be constructed and/or obtained in any suitable way, as the manner in which the knowledge representation is constructed and/or obtained is not a limitation of aspects of the described methods and systems.

[0070] Regardless of which knowledge representation is used in decision block 302, the determination of whether the relevant portion of the user context information matches a concept in the knowledge representation may be made in any suitable way. In some embodiments, the relevant portion of the user context information may be compared with a concept identifier. For example, when the relevant portion of the user context information is represented by an alphanumeric string, the alphanumeric string may be compared with a string identifying the

concept (sometimes referred to as a "concept label") to determine whether or not there is a match. The match may be an exact match between the strings, or a substantially exact match in which all words, with the exception of a particular set of words (e.g., words such as "and," "the," "of," etc.), must be matched. Moreover, in some embodiments, the order of words in the strings may be ignored. For instance, it may be determined that the string "The Board of Directors," matches the concept label "Board Directors" as well as the concept label "Directors Board."

[0071] If it is determined, in decision block 302, that the relevant portion of the user context information matches a concept in the knowledge representation, process 300 proceeds to decision block 304, where it is determined whether there are multiple concepts in the knowledge representation matching the relevant portion. For example, the selected portion of the user context information may be an alphanumeric string "bark" indicating that the user may be interested in information about "bark." However, it may not be clear whether the user is interested in information about the bark of a dog or the bark of a tree; there may be concepts associated to each such concept in the knowledge representation.

[0072] If it is determined, in decision block 304, that there is only one concept, in the knowledge representation matching the relevant portion of the user context information, the one concept is identified as the active concept and process 300 proceeds via the NO branch to act 320 where the active process is returned for subsequent processing, for example, as described in greater detail below with reference to acts 106-110 of process 100.

[0073] On the other hand, if it is determined that there are multiple concepts in the knowledge representation matching the relevant portion of the user context information, process 300 continues via the YES branch to acts 306-308, where one of the matching concepts may be selected as the active concept. This may be done in any suitable way. In some embodiments, one of the multiple matching concepts may be selected by using a disambiguation process.

[0074] Any suitable disambiguation process may be employed to identify an active concept among the multiple concepts matching the relevant portion of the user context information. Such a disambiguation process may comprise using one or more disambiguation terms to identify the active concept among the multiple concepts such that the identified active concept is likely to represent information that the user may be interested in. The disambiguation

process may comprise generating a set of candidate disambiguation terms and selecting one or more candidate disambiguation terms to use for identifying the active concept. For example, a set of candidate disambiguation terms, including the terms "dog" and "tree," may be generated. Subsequent selection of the disambiguation term "dog", which may be performed either automatically or based at least in part on user input, may indicate that the user is interested in information about "dog barking." As such, the selected disambiguation terms may be used for semantically disambiguating among the multiple concepts identified in act 304 to identify the active concept.

[0075] Accordingly, in act 306, a set of candidate disambiguation terms may be generated. This may be done in any suitable way. For example, the set of candidate disambiguation terms may comprise one or more keywords, morphemes, and/or any other suitable knowledge representation entities of one or more concepts among the multiple concepts matching the relevant portion of the user context information. Additionally, the set of candidate disambiguation terms may comprise one or more keywords, morphemes, and/or any other suitable KR entities of any concepts connected, within a predetermined degree of separation in the semantic network, to a concept among the multiple concepts. Any suitable degree of separation (e.g., one, two, three, four, five, etc.) may be used. In some embodiments, the set of candidate disambiguation terms may not comprise any of the terms in the relevant portion of the user context information, though in other embodiments, the set of candidate disambiguation terms may comprise one or more terms in the relevant portion of the user context information.

[0076] Next, process 300 proceeds to act 308, where one or more of the candidate disambiguation terms may be selected. The selection may be performed in any suitable way and may be performed automatically and/or may involve obtaining one or more disambiguation terms based on user input. For example, in some embodiments, one or more candidate disambiguation terms may be provided to the user, such that the user may select those terms that are indicative of what the user is interested in. The candidate disambiguation terms may be provided to the user in any of a variety of possible ways. For example, in some embodiments, the terms may be provided from server 208 (i.e., the computer that performs process 100) to the application program 204 on client 202 from which the user context information may have been obtained. In embodiments in which application program 204 is an intranet or WWW browser, the

terms may be provided in the form of a web page. As such, the user may select one or more terms to indicate the type of information that the user may be interested in.

[0077] Regardless of the manner in which one or more candidate disambiguation terms may be provided to a user, user input comprising a selection of one or more disambiguation terms may be received as part of act 308 of process 300. For example, application program 204 that received the set of candidate disambiguation terms generated in act 306 may accept input from the user selecting one or more of the terms, and may send an indication of the user-selected term(s) to the server executing process 100.

[0078] In some embodiments, one or more disambiguation terms may be selected automatically from the set of candidate disambiguation terms, without requiring any user input. For example, one or more terms from the set of candidate disambiguation terms, generated in act 306, may be selected based on user context information (e.g., the user's browsing history, online profile, user selected preferences, or any other type of user context information described earlier). Consider, for example, a situation in which a user is searching for "bark," but that it is clear from the user's browsing history that the user has shown interest in various information about dogs. In this case, it is likely that the user is searching for information about a "dog bark" rather than "tree bark." Accordingly, the user context information may be used to select the term "dog" from the set of candidate disambiguation terms {"dog" and "tree"}. As another example, the user's online profile on a social networking website may indicate that the user is an avid botanist (or geo-spatial information associated with the user indicates that the user is located in a rainforest), in which case it is likely that the user is searching for information about "tree bark" rather than "dog bark." Though it should be recognized that the above described techniques for selecting disambiguation terms are merely illustrative as the disambiguation terms may be selected in any other suitable manner.

[0079] Regardless of the manner in which one or more disambiguation terms may be obtained, in act 308, the obtained terms may be used to identify an active concept among the multiple concepts matching the relevant portion of the user context information. Accordingly, the identified active concept may represent information in which one or more users may be

interested. After the active concept is identified, in act 308, process 300 proceeds to act 320 where the active process is returned for subsequent processing and process 300 completes.

[0080] Consider, again, decision block 302 of process 300. If it is determined in decision block 302 that the relevant portion of the user context information does not match any concept in the knowledge representation (the NO output branch), process 300 proceeds to act 310, where the relevant portion of the user context information may be decomposed into one or more knowledge representation entities. For example, the relevant portion of the user context information may be decomposed into individual concepts, keywords, and/or morphemes. This may be done in any suitable way. For example, when the portion of the user context information is represented by an alphanumeric string, the string may be tokenized or separated into more elemental knowledge representation entities. Stop words such as "the" and "and" may be filtered out or ignored. For example, if the alphanumeric string is a user's search query "The BP Board of Directors," the query may be tokenized into the following entities: "Board of Directors," "BP Board," "BP Directors," "BP", "Board," and "Directors." It should be recognized that many other techniques may be applied to separating the relevant portion of the user context information into knowledge representation entities including the semantic analysis methods described in U.S. Patent Application Serial No. 13/165,423, filed 06/21/2011, and titled "Systems and Methods for Analyzing and Synthesizing Complex Knowledge Representations," which is incorporated herein by reference.

[0081] Process 300 continues to act 312, where concepts in the knowledge representation that cover the KR entities, which were derived in act 310, are identified. This may be done in any suitable way and, for example, may be done by comparing each of the KR entities with concepts in the KR to see if there is a match. In some embodiments, a string associated with a KR entity may be compared with labels of concepts in the KR. For example, consider semantic network 401 shown in FIG. 4A comprising concepts 402 and 406 labeled "Board of Directors" and "Board," respectively. Concepts 402 and 406 are connected by a "defined-by" edge 404. Though not explicitly shown, the node associated with the concept labeled "Board of Directors" may also be connected via a "defined-by" edge to a node associated with the concept labeled "Director." Accordingly, nodes existing in semantic network 401 cover KR entities "Board of Directors," "Directors," and "Board." Note that these KR entities were derived from the alphanumeric string

"BP Board of Directors" in act 310 of process 300. Note that semantic network 401 does not include a concept with the label "BP Board of Directors."

[0082] Process 300 next continues to acts 314-318, where an active concept corresponding to the relevant portion of the user context information may be generated and, subsequently, used to construct a user-specific knowledge representation comprising the active concept. First, in act 314, an active concept associated with the relevant portion of the user context information may be generated. This may be done in any suitable way. For example, a new node may be added to the knowledge representation and may be associated with the relevant portion of the user context information. As such, the node may be assigned an identifier (i.e., a concept label) comprising the relevant portion of the user context information. For example, as shown in FIG. 4B, node 408 associated with the generated active concept and labeled "BP Board of Directors" was added to semantic network 401 to form semantic network 403. In this and other diagrams of semantic networks, the node corresponding to an active concept may be indicated by a rectangle.

[0083] Next, as part of act 316, the new node may be connected by one or more new edges to one or more concepts already present in the knowledge representation. The new node, representing the generated active concept, may be connected to any suitable concepts in the knowledge representation and, for example, may be connected to one or more concepts in the knowledge representation that cover the knowledge representation entities derived from the relevant portion of the user context information. Thus, in the "BP Board of Directors" example, node 408 may be connected to the node associated with the concept "Board of Directors," to the node associated with the concept "Board," and/or to the node associated with the concept "Directors."

[0084] In some embodiments, the new node may be connected to nodes associated with the most complex concepts that cover the KR entities derived in act 310. Complexity of a concept may be defined in any of numerous ways. For example, complexity of a concept may be indicative of the number of other concepts that are "defined-by" the concept; the greater the number of concepts "defined-by" the concept, the greater its complexity. Thus, complexity of a concept may be proportional to the number of outgoing "defined-by" edges from the node

corresponding to that concept. In semantic network 401, for example, the concept "Board of Directors" has a greater complexity than the concept "Board." As another example, complexity of a concept may be indicative of the number of words in the label of the concept; the greater the number of words, the greater its complexity. Thus, complexity of a concept may be proportional to the number of words in the concept label. In this case, the concept "Board of Directors" also has a greater complexity than the concept "Board." Accordingly, node 408, associated with the concept "BP Board of Directors," is connected, in semantic network 405, by a new "defined-by" edge 410 to node 402, corresponding to the "Board of Directors" concept.

[0085] Finally, in act 318, the knowledge representation may be further augmented such that the knowledge representation includes concepts that may cover all of the KR entities derived from the relevant portion of the user context information in act 310. To this end, a new node may be added to the knowledge representation for each KR entity derived in act 310 but not covered by a concept already in the knowledge representation. Each such new node may be connected to one or more concepts existing in the knowledge representation and, for example, may be connected via a "defined-by" edge to the concept associated with the generated active concept. For example, the keyword "BP" was derived from the new "BP Board of Directors" concept, but is not covered by any of the concepts in semantic networks 401 or 403. Thus, as shown in FIG. 4C, node 416 associated with the concept "BP" may be added to the knowledge representation in act 318 and may be connected to node 408, associated with the active concept, via a "defined-by" edge 414.

[0086] Thus, a new user-specific knowledge representation is created after acts 314-318 have been executed. The created knowledge representation is user-specific because it comprises one or more concepts derived from user context information associated with one or more users and the knowledge representation used in act 302. In the examples of FIGS. 4A-4C, semantic network 405 was created by incorporating two concepts (i.e., "BP" and "BP Board of Directors" obtained from user context information) into semantic network 401. Though, it should be recognized that the examples of FIGS. 4A-4C are merely illustrative and are not limiting on aspects of the present invention. Next, process 300 continues to act 320, where the active concept generated in acts 314-318 may be provided for subsequent processing, and after act 320, process 300 completes.

[0087] It should be appreciated that after the user-specific knowledge representation is created, it may be stored such that it may be subsequently used in any processing associated with the one or more users. For example, the user-specific knowledge representation may be used and/or updated anytime processes 100 or 300 may be executed in connection with the user(s). The user-specific knowledge representation may be stored in any suitable way and, for example, may be stored using one or more non-transitory, tangible computer-readable storage media of any suitable type.

III. Identify Concepts Relevant to Active Concept

[0088] As discussed above, at act 106 of process 100, one or more concepts relevant to the active concept may be obtained by using a knowledge representation comprising the active concept. The knowledge representation may be any suitable knowledge representation and, in some instances, may be a user-specific knowledge representation associated with the user(s) whose context information was obtained in act 102 of process 100. The active concept may be any suitable concept in the knowledge representation and may be identified based at least in part on the user context information, for instance, by using process 300, or in any other suitable way.

[0089] In some embodiments, one or more concepts relevant to the active concept may be obtained based at least in part on the structure of the knowledge representation comprising the active concept. For example, when the knowledge representation is a semantic network, one or more concepts relevant to the active concept may be obtained based at least in part on the structure of the graph that represents the semantic network. Any of numerous aspects of the graph structure may be used including, but not limited to, the directionality of the edges representing semantic relationships between concepts and whether the semantic relationships are "defined-by" relationships or "is-a" relationships. Additionally or alternatively, the structure of the graph may be used to synthesize one or more new concepts, not initially in the semantic network, that may be relevant to the active concept. In this case, any of the synthesized concepts may be used to augment the semantic network.

[0090] Any of numerous techniques for obtaining concepts relevant to the active concept based on the graph structure of the semantic network comprising the active concept may be used in act 106 of process 100. In some embodiments, concepts relevant to the active concept may be

obtained by performing one or more types of operations with respect to the graph structure of the semantic network. Three such operations, namely: (1) retrieval, (2) addition, and (3) substitution, are described in greater detail below. These three operations are merely illustrative, however, and any other suitable operations for identifying concepts relevant to the active concept, based at least in part on the graph structure of the semantic network, may be used. For brevity, in the description of the operations that follows, no explicit distinction is made between a node in a graph used to represent a concept and the concept itself. Thus, an edge between two concepts corresponds to an edge between the nodes in the semantic graph used to represent those two concepts.

[0091] A retrieval operation may be used to identify concepts in the semantic network that are relevant to the active concept. In some embodiments, the retrieval operation may be used to identify concepts that were represented in the semantic network before the active concept was identified and/or generated. Though, in other embodiments, the retrieval operation may be used to identify concepts that were added to the semantic network when the active concept was generated (e.g., in act 318 of process 300).

[0092] In some embodiments, the retrieval operation may identify a concept that is connected by one or more edges, of any suitable type and/or direction, to the active concept as a concept relevant to the active concept. For example, the retrieval operation may identify a concept that is connected by one or more "is-a" edges to the active concept as a concept relevant to the active concept. As a specific example, the retrieval operation may identify a concept that has outgoing "is-a" edge towards the active concept or a concept that has an incoming "is-a" edge from the active concept as a concept relevant to the active concept.

[0093] A simple example of a retrieval operation is illustrated in FIG. 5A, which shows a semantic network comprising the active concept "press" (that the concept "press" is active is indicated here by a rectangle) and another concept "push press." The concept "push press" is connected to the active concept via an outgoing "is-a" relationship. Accordingly, a retrieval operation may be used to identify the concept "push press" as a concept relevant to the active concept. Note that in FIGS. 5A-5H, the concepts identified as concepts relevant to the active concept are indicated by a diamond.

[0094] In contrast to the retrieval operation, which may be used to obtain concepts relevant to the active concept among the concepts already in the semantic network, the addition and substitution operations described below may be used to obtain concepts relevant to the active concept by synthesizing new concepts based at least in part on the active concept and on the concepts in the semantic network. Note that in FIGS. 5B-5H, concepts added to the active concept to synthesize a new concept are indicated by a hexagon.

[0095] An addition operation may synthesize a new concept by using the active concept and at least one other concept in the semantic network, and return the new concept as a concept relevant to the active concept. The new concept may be synthesized using any of numerous techniques including at least: (1) attribute co-definition, (2) analogy-by-parent, (3) analogy-by-sibling, (4) attribute commonality or any suitable combination thereof.

[0096] In some embodiments, an addition operation may be used to synthesize a new concept by using the attribute co-definition technique. A first concept in a semantic network is an attribute of a second concept in the semantic network if the first concept defines the second concept. This may be ascertained from the graph representing the semantic network if there is an outgoing "defined-by" edge from the second concept to the first concept. For example, as shown in FIG. 5B, the concepts "bench," "press," and "sets" are attributes of the concept "bench press sets," and the concepts "press" and "sets" are attributes of the concept "press sets." If the active concept is an attribute of (i.e., is connected via an outgoing "defined by" edge to) a first concept, and the first concept has one or more other concepts as an attribute, it may be said that the active concept and the other concept(s) co-define the first concept. For example, in FIG. 5B, "press" is the active concept and, the concepts "press," "bench," and "sets" co-define the concept "bench press sets."

[0097] In the attribute co-definition technique, a new concept may be synthesized by combining the active concept with any of the other concepts co-defining a concept with the active concept. For example, as shown in FIG. 5B, the concept "press sets" may be synthesized by combining "press" and "sets." As another example (not shown in FIG. 5B), the concept "bench press" may be synthesized by combining "press" and "bench."

[0098] In some embodiments, an addition operation may be used to synthesize a new concept by using the analogy-by-parent technique. In a semantic network, a first concept with an outgoing "is-a" edge to a second concept may be considered as a child concept of the second concept. Stated differently, the second concept may be considered to be a parent concept of the first concept. For example, in FIG. 5C, the concept "seat" is a parent concept of the active concept "recliner." The analogy-by-parent technique is motivated by the idea that an attribute that co-defines a concept with a parent of the active concept may be relevant to the active concept. Stated more plainly, something relevant to the parent may be relevant to the child.

[0099] Accordingly, in the analogy-by-parent technique, a new concept may be synthesized by using the active concept and any second concept that, together with the parent of the active concept, co-defines (or partially co-defines) a third concept. For example, in FIG. 5C, the concept "seat," which is the parent concept of "recliner," and "toilet" together co-define the concept "toilet seat." Thus, the concept "toilet" co-defines another concept with and, as such, may be deemed relevant to a parent of the active concept. Accordingly, the active concept "recliner" and the concept "toilet" may be used to synthesize a new concept "recliner toilet."

[00100] In some embodiments, an addition operation may be used to synthesize a new concept by using the analogy-by-sibling technique. In a semantic network, any two concepts with outgoing "is-a" edges ending at a common (parent) concept may be considered siblings of one another. For example, in FIG. 5D, the concepts "chair" and "recliner" may be considered as siblings. The analogy-by-parent technique is motivated by the idea that an attribute that co-defines a concept with a sibling of the active concept may be relevant to the active concept. Stated more plainly, something relevant to one sibling may be relevant to another sibling.

[00101] Accordingly, in the analogy-by-parent technique, a new concept may be synthesized by using the active concept and any second concept that, together with the sibling of the active concept, co-defines (or partially co-defines) a third concept. For example, in FIG. 5D, the concept "chair," which is a sibling of the active concept "recliner," and "massage" together co-define the concept "massage chair." Thus, the concept "massage" co-defines another concept with and, as such, may be deemed relevant to, the sibling concept "chair." Accordingly, the

active concept "recliner" and the concept "massage" may be used to synthesize a new concept "massage recliner."

[00102] It should be recognized that the terms "parents" and "siblings" are used to provide intuition behind some of the above-described operations and that, in some embodiments, such as an atomic knowledge representation model, concepts may not include literal "parent" and "sibling" relationships. The terms "siblings" and "parents" suggest a taxonomy structure in a complex knowledge representation. In contrast, in some embodiments, an atomic knowledge representation model may include only "is-a" and "defined-by" relations.

[00103] In some embodiments, an addition operation may be used to synthesize a new concept by using the attribute commonality technique. In a semantic network, two concepts may be said to exhibit "attribute commonality" if the concepts share one or more attributes with one another. For example, as shown in FIG. 5E, the concept "massage chair" and "shiatsu therapy massage chair" share the attributes "massage" and "chair," and, as such, may be said to exhibit attribute commonality. The attribute commonality technique is motivated by the idea that if a first concept shares one or more attributes with a second concept, then any other attributes of the second concept may be relevant to the first concept.

[00104] Accordingly, in the attribute commonality technique, a new concept may be synthesized by using the active concept and any attribute of a second concept that shares one or more attributes with the active concept. For example, as shown in FIG. 5E, the active concept "massage chair" and "shiatsu," which is an attribute of the concept "shiatsu therapy massage chair" that has attribute commonality with "massage chair," may be used to synthesize a new concept "shiatsu massage chair." As another example, not shown in FIG. 5E, the active concept "massage chair" and "therapy," may be used to synthesize a new concept "massage therapy chair."

[00105] In some embodiments, the attribute commonality technique may comprise generating a new concept by using the active concept and another concept identified by using at least one "is-a" bridge. In a semantic network, two concepts are connected via an "is-a" bridge if they both share outgoing "is-a" edges terminating at a common (parent) concept. For example, in FIG. 5F, the concepts "yoga" and "weightlifting" are connected via an "is-a" bridge to the

concept "exercise." Also, the concepts "mat" and "bench" are connected via an "is-a" bridge to the concept "furniture."

[00106] In the attribute commonality technique, a new concept may be synthesized by using the active concept and a second concept that has an attribute connected to an attribute of the active concept via an "is-a" bridge. In some instances, the new concept may be synthesized by using the active concept and any attribute of the second concept that is not connected to an attribute of the active concept via an "is-a" bridge. For example, as shown in FIG. 5F, the active concept "yoga mat" has "yoga" and "mat" as its attributes. Each of these attributes is connected with an attribute of the concept "sweat-absorbent weightlifting bench." Accordingly, the active concept "yoga mat" and the attribute "sweat-absorbent" may be used to synthesize a new concept "sweat-absorbent yoga mat."

[00107] The above-described techniques for performing an addition operation comprise synthesizing a new concept, as a concept that may be relevant to the active concept, by combining the active concept with another concept. As a result, the synthesized concept may be a less general or "narrower" concept than the active concept. However, it should be recognized that, concepts relevant to the active concept need not be less general than the active concept and, indeed, may be more general or "broader" concepts than the active concept.

[00108] Accordingly, in some embodiments, one or more attributes of the active concept may be pruned in order to produce a candidate that is more general than the active concept. This may be done in any suitable way. For instance, attributes may be pruned by performing an "inverse" addition operation, wherein an attribute of the active concept may be removed if, according to any of the above-described techniques, that attribute may be combined with the "broader" concept that results from the pruning. For example, if in the semantic network shown in FIG. 5E, the concept "shiatsu massage chair" was an active concept and the concept "massage chair" was not in the network, then the concept "massage chair" may be created by pruning the attribute "shiatsu." Although, in some embodiments, any suitable attribute may be pruned so long as the resulting concept is not in the semantic network. For example, the attribute "massage" may be pruned resulting in the concept "shiatsu chair."

[00109] Another operation that may be used to obtain one or more concepts relevant to the active concept is the substitution operation. The substitution operation may be used to synthesize a new concept by replacing one or more attributes of the active concept with another concept, which may be a broader or a narrower concept than the attribute that it replaces. To perform a substitution, either a retrieval or an addition operation may be performed on one or more attributes of the active concept. The concept identified or generated by the retrieval or addition operations, respectively, when performed on a specific attribute of the active concept, may be used to replace the specific attribute to synthesize a new concept.

[00110] Consider, for example, the semantic network shown in FIG. 5G comprising the active concept "strict press" having attributes "strict" and "press." As previously described with reference to FIG. 5A, a retrieval operation performed on the attribute "press" may be used to identify the (narrower) concept "push press." According to the substitution technique, this narrower concept ("push press") may be combined with any attribute or attributes of the active concept (other than the attribute from which the narrower concept was derived) to synthesize a new concept. In this way, the concept "strict push press" may be synthesized. Similarly, a substitution with retrieval operation may be performed to substitute an attribute of the active concept with a concept that is broader than that attribute.

[00111] FIG. 5H illustrates performing a substitution operation by using an addition operation based on the attribute co-definition technique. In this example, applying the attribute co-definition technique to the attribute "press" of the active concept "push press" results in the concept "press sets," as previously described with reference to FIG. 5B. Accordingly, the concept "press sets" may be used to replace the attribute "press" in the concept "push press" to synthesize a new concept "push press sets."

[00112] In the same vein, substitution operations using any other type of addition operation (e.g., analogy-by-parent, analogy-by-sibling, and attribute commonality) on one or more attributes of the active concept may be used to synthesize one or more concepts relevant to the active concept.

IV. Score and Select Identified Concept(s) by using Relevance Measure or Measures

[00113] After one or more concepts relevant to the active concept are obtained in act 106, process 100 proceeds to act 108, where the obtained concepts may be scored and a subset of the concepts may be selected for subsequent use based on the calculated scores. Scores associated to the concepts obtained in act 106 may be calculated in any of numerous ways. In some embodiments, the scores may be obtained by using one or more relevance measures indicative of how relevant a concept to be scored may be to the active concept. A relevance measure may be computed based at least in part on the structure of the graph that represents the semantic network containing the concept to be scored and the active concept.

[00114] Five measures of relevance are described in greater detail below along with some of their variations, namely: (1) generation certainty, (2) concept productivity, (3) Jaccard (4) statistical coherence, and (5) cosine similarity. Though it should be recognized that these techniques are merely illustrative and that any other suitable techniques for assigning a score to a concept may be used. For example, as described in greater detail below, any of the above techniques may be combined to calculate an integrated score for a concept obtained in act 106.

IV.A Generation Certainty Technique

[00115] In the generation certainty technique, concept scores may be calculated based at least in part on the structure of the semantic network comprising the concepts. Recall that any of the concepts obtained in act 106 of process 100 may be in the semantic network or may be added to the semantic network after they are synthesized. The generation certainty score calculated for a particular concept may depend on the structure of the semantic network as well as the locations of the particular concept and the active concept within the semantic network. The score may depend on any of numerous aspects of the structure of the semantic network including, but not limited to, the number of edges in a path between the active concept and the particular concept, the number of nodes in a path between the active concept and the particular concept, the types of edges in a path between the active concept and the particular concept, the types of nodes in a path between the active concept and the particular concept, the directionality of the edges in a path between the active concept and the particular concept, any weights associated with edges in a path between the active concept and the particular concept, and any suitable combination thereof. It should be recognized that the structure of a graph representing the semantic network

may be such that there are one or multiple paths between the active concept and a concept to be scored.

[00116] In some embodiments, for example, the generation certainty score computed for a particular concept may be inversely proportional to the number of edges and/or nodes separating the active concept and the particular concept in the semantic graph. Accordingly, the score computed for a concept separated by a large number of edges and/or nodes from the active concept may be lower than the score computed for a concept separated by a smaller number of edges and/or nodes from the active concept.

[00117] As previously mentioned, in some embodiments, the generation certainty score may be calculated as a function of the weights associated with edges in the semantic network. In particular, the generation certainty score may be calculated as a function of the weights associated with a set of edges in a path between the active concept and the concept being scored. In this case, the generation certainty score may be calculated by taking a product of the weights of the edges in the path from the active concept to the concept being scored.

[00118] A weight may be assigned to an edge in a semantic network in any of numerous ways. In some embodiments, the weight assigned to an edge may be computed based on a measure of certainty associated with traversing that edge. In turn, the amount of certainty associated with traversing an edge may depend on the type of the edge (i.e., is the edge a "defined-by" edge or an "is-a" edge) and/or on the direction of traversal. In some embodiments, the weight assigned to an edge may be a number between 0 and 1, but in other embodiments the weight may be a number in any other suitable range.

[00119] For example, traversal of a "defined-by" edge may reduce the certainty associated with traversing the edge by a factor of x, where x may be any suitable number between 0 and 1 and, for example, may be any factor greater than or equal to 0.25, 0.5, 0.75, 0.9, etc. Similarly, traversal of an "is-a" edge may reduce the certainty of traversing the edge by a factor of y, where y may be any suitable number between 0 and 1 and, for example may be any factor greater than or equal to 0.25, 0.5, 0.75, 0.9, etc. In some instances, the factor x may be equal to the factor y, but, in some instances, these factors may be different such that the amount of certainty associated with traversing an edge may depend on the type of edge being traversed.

[00120] In some embodiments, the amount of certainty associated with traversing an edge may depend on the directionality of the edge and the direction that the edge may be traversed when traversing that edge in a path from one concept to another. For instance, traveling from concept A to concept B (where A "is-a" B such that there is an outgoing "is-a" edge from the node associated with concept A to the node associated with concept B in the semantic network) may reduce the amount of certainty by one factor (e.g., 0.9) while traveling against the direction of the "is-a" edge may reduce certainty by a different factor (e.g., 0.8).

[00121] In some embodiments, the generation certainty score assigned to a concept may depend on whether that concept was obtained by using a retrieval operation (i.e., the concept was already in the semantic network) or obtained by using an addition or substitution operation (i.e., the concept was synthesized). For example, the generation certainty score may be reduced by a factor (e.g., 0.25) when the concept was synthesized.

[00122] One illustrative, non-limiting example of computing a generation certainty score is shown in FIG. 6A. In this case, the generation certainty of the candidate "massage recliner" may be calculated as a product of the weights associated with edges in the path from the active concept "recliner" to the synthesized concept "massage recliner." As shown in FIG. 6A, the weight associated with each of the "defined-by" edges along the path is 0.9 and the weight associated with each of the "is-a" edges along the path is 0.75. Further, since the concept "massage recliner" is a synthesized concept, the overall generating certainty score is adjusted by a factor of 0.25. Thus, the generation certainty score, S_{gc}, may be calculated according to:

$$S_{gc} = Edge_{recliner\text{-}seat} \times Edge_{seat\text{-}chair} \times Edge_{chair\text{-}massage\ chair} \times Edge_{massage\ chair\text{-}massage} \times Node_{massage\ recliner}$$

$$= Edge_{is\text{-}a} \times Edge_{is\text{-}a} \times Edge_{defined\text{-}by} \times Edge_{defined\text{-}by} \times Node_{synthesized}$$

$$= 0.9 \times 0.9 \times 0.75 \times 0.75 \times 0.25 = 0.1139$$

[00123] It should also be recognized that, in some embodiments, the numerical values of the weights associated with the edges in a semantic network may be manually assigned (e.g. assigning the weight of 0.9 to an "is-a" edge and a weight of 0.75 to a "defined-by" edge). Additionally or alternatively, the numerical values of the weights may be based on the statistical

coherence measures described below and/or calculated as probabilities using the teachings disclosed in U.S. Provisional Application Ser. No. 61/430,810, filed on January 7, 2011, titled "Probabilistic Approach for Synthesis of a Semantic Network"; U.S. Provisional Application Ser. No. 61/430,836, filed on January 7, 2011, titled "Constructing Knowledge Representations Using Atomic Semantics and Probabilistic Model"; and U.S. Provisional Application Ser. No. 61/532,330, filed on September 8, 2011, titled "Systems and Methods for Incorporating User Models and Preferences into Analysis and Synthesis of Complex Knowledge Representation Models," all of which are hereby incorporated by reference in their entireties.

IV.B Concept Productivity Score

[00124] In the concept productivity technique, the score of a concept may be calculated based on the number of other concepts in the semantic network that the concept defines. For example, the concept productivity score of a concept may be calculated based on the number of incoming "defined-by" edges that the concept possesses. Some further examples are provided below.

[00125] For example, the concept productivity score assigned to a concept obtained by using a retrieval operation (e.g., as described with reference to act 106 of process 100), may be calculated based on the number of concepts the concept defines. For example, the active concept may be "press", which may have an incoming "is-a" relationship with the concept "push press" and an incoming "is-a" relationship with "dumbbell press." As such, both of these concepts may be retrieved as concepts relevant to the active concept in act 106. However, if the number of concepts defined by the concept "push press" is greater than the number of concepts defined by the concept "dumbbell press", then the concept "push press" will be assigned a higher concept productivity score than the concept "dumbbell press."

[00126] As another example, the concept productivity score assigned to a concept obtained by using an addition operation (e.g., as described with reference to act 106 of process 100), may be calculated based on the number of concepts defined by the concept to be added to the active concept in order to generate the synthesized concept. For example, the active concept may be "press" and the concepts synthesized by using one of the addition operations may be "press sets" or "press movements." If the number of concepts defined by "movements" is greater

than the number of concepts defined by "sets," then the concept "press movements" will be assigned a higher concept productivity score than the concept "press sets."

[00127] As another example, the concept productivity score assigned to a concept obtained by using a substitution operation may be calculated based on the number of concepts defined by the concept to be substituted for one of the attributes of the active concept. For example, the active concept may be "push press" and its attribute "press" may be substituted with the concept "press sets" or the concept "press movements." If the number of concepts defined by the concept "press movements" is greater than the number of concepts defined by the concept "press sets," then the synthesized concept "push press movements" will be assigned a higher concept productivity score than the synthesized concept "push press sets."

IV.C Jaccard Score

[00128] In the Jaccard score technique, the score of a particular concept may be calculated based on the number of concepts that fall within a particular degree of separation from the active concept as well as within the same degree of separation from the particular concept. For example, when the degree of separation is one, the Jaccard score of a particular concept may be calculated based on the number of neighbors in common between the particular concept and the active concept. In a semantic network, the neighbor of concept A may be any concept sharing any type of edge with concept A. Since, they share an edge, a neighbor of a concept is within a single degree of separation from the concept. In this case, the larger the number of neighbors in common between an active concept and a concept to be scored, the higher the Jaccard score assigned to that concept. As such, the Jaccard score provides an indication of how interconnected two concepts may be in the semantic network, which, in turn, may be an indication of the relevance of the two concepts. Though, it should be recognized that any degree of separation (e.g., one, two, three, four, five, six, seven, etc.) may be employed.

[00129] A Jaccard index is a similarity measure for measuring similarity between two sets A and B. In some instances, the Jaccard index may be defined as the size of the intersection of sets A and B divided by the size of the union of the sets A and B, as shown below:

$$J(A,B) = \frac{|A \cap B|}{|A \cup B|}.$$

[00130] The Jaccard index may be applied in our case as follows. Let the set A represent the set of concepts that may be neighbors, or may be within a predetermined number of degrees of separation in the semantic network, from the active concept. Let the set B represent the set of concepts that may be neighbors, or may be within a predetermined number of degrees of separation in the semantic network, from the concept to be scored. Thus, the denominator in the above equation represents the total number of concepts that may be neighbors (or may be within a predetermined number of degrees of separation) of the active concept and/or the concept to be scored while the numerator represents the total number of concepts that are both a neighbor (or may be within a predetermined number of degrees of separation) of the active concept and the concept under evaluation. Accordingly, the Jaccard score of a concept may be computed as the Jaccard index.

[00131] An example of computing a Jaccard score for a concept obtained in act 106 of process 100 is shown in FIG. 6B. In this example, the neighborhood of a concept has been selected as comprising concepts within two degrees of separation of the concept. Though, it should be recognized, that a neighborhood associated with any suitable degree of separation may be used. Accordingly, all concepts within two degrees of either the active concept or concept 602, which is the concept to be scored, have been indicated with diagonal lines, unless they are within two degrees of separation from both the active concept and concept 602; such concepts, being within two degrees of both the active concept and concept 602, are indicated with vertical lines. To compute the Jaccard score, observe that the number of concepts with either diagonal or vertical lines is the number of concepts (25) in the denominator of the Jaccard score. The number of concepts with vertical lines (7) is the number of concepts in the numerator of the Jaccard score. Accordingly, the Jaccard score of concept 602 would be calculated as 7 divided by 25, or 0.38.

[00132] In some embodiments, such as the embodiment illustrated in FIG. 6B, neither the active concept nor the concept to be scored are considered neighbors of themselves or of one another. However, in other embodiments, the active concept and/or the concept to be scored may be considered as neighbors of themselves and/or of one another. In the illustration of FIG. 6B, for example, if the active concept and concept 602 were to be considered neighbors of

themselves and of one another, the Jaccard score would be computed as 9 divided by 27, or 0.333.

[00133] In some embodiments, the Jaccard score may be calculated as the complement of the Jaccard index according to 1- J(A,B) such that the Jaccard score may be indicative of a measure of dissimilarity between concepts A and B. It should be recognized that in this case, concepts with lower scores (rather than higher scores) may be selected in act 108 of process 100. Further, in this case the concept with a higher Jaccard score would then be considered to possess a weaker relationship with the active concept than a concept with a lower score.

[00134] In cases when a Jaccard score is being computed for a concept obtained by using a retrieval operation (e.g., as described with reference to act 106), the Jaccard score may be obtained by applying the above-described techniques to the retrieved concept. However, when a Jaccard score is being computed for a concept synthesized via an addition operation, the Jaccard score may be obtained by applying the above-described techniques not to the synthesized concept, but rather to the concept that was combined with the active concept to produce the synthesized concept (e.g., the concept "shiatsu" shown in FIG. 5E). Similarly, when a Jaccard score is being computed for a concept synthesized via a substitution operation, the Jaccard score may be obtained by applying the above-described techniques not to the synthesized concept, but to the concept that was used to substitute an attribute of the active concept as part of the substitution (e.g., the concept "press sets" shown in FIG. 5H).

IV.D Statistical Coherence Score

[00135] Another technique for computing scores for concepts obtained in act 106 is the so-called "statistical coherence" technique where the statistical coherence score assigned to a particular concept will depend on the frequency of co-occurrence of that concept with the active context in one or more text corpora. As such, a concept that co-occurs more frequently with the active concept may be more relevant to the active concept than a concept that co-occurs less frequently with the active concept.

[00136] Any suitable corpus or corpora may be used to calculate the statistical coherence score for a concept obtained in act 106 of process 100. For example, the corpora may be from a

single source (e.g., all content found at URLs containing the string "wikipedia.org") or multiple sources. As another example, subject-specific corpora may be used such as corpora containing content about politics, medical articles, sports, etc. Each corpus may be of any suitable type and, for example, may be a text corpus or a corpus containing multiple types of content.

[00137] Regardless of the number and types of corpora used for calculating coherence scores, in some embodiments, the active concept may be used to select a subset of content in the corpora (e.g., a portion of the documents in a text corpus) to use for calculating statistical coherence scores. This may be done in any suitable way. For example, the active concept may be used to select content relevant to the active concept, which, for example, may be only that content which contains the label of the active concept.

[00138] In some embodiments, the content used for statistical coherence calculations may be further restricted to that content which contains the active concept and at least one of the concepts in a neighborhood of the active concept in the semantic network. Recall that such a neighborhood may include all concepts, in the semantic network, that are within a predetermined number of degrees of separation (e.g., one, two, three, four, etc.) from the active concept. The additional restriction may be accomplished in any suitable way and, for example, may be achieved by using only that content which contains the label of the active concept and the label of at least one other concept in the neighborhood of the active concept. Restricting the corpora based on the active concept as well as its neighbors may be advantageous in that content that includes the label of the active concept, but is directed towards a distinct meaning, is not considered when calculating a statistical coherence score.

[00139] For example, if the active concept is "bat" and the concepts found within the active concept's neighborhood include the concepts "baseball bat," "club," "paddle," and "lumber," the content used for calculating statistical coherence scores may be limited to content that includes the active concept "bat" and at least one of the neighboring concepts "baseball bat," "club," "paddle," and "lumber." The inclusion of at least one of these neighbors may increase the likelihood that documents that include the concept "bat," but are related to the mammal, are avoided when calculating the statistical coherence score.

[00140] Accordingly, in some embodiments, the statistical coherence score may be computed as a function of the ratio of the number of documents containing both the active concept, at least one concept in a neighborhood of the active concept, and the concept to be scored and the number of documents containing the active concept and the at least one concept in the neighborhood of the active concept. The function may be any suitable function and, for example, may be the identity function or any other suitable monotonically increasing function (e.g., logarithm). When calculated in this manner, the statistical coherence score may reflect the proportion of the total number of documents relevant to the active concept that are also relevant to the concept to be scored. Accordingly, the higher the statistical coherence score for a concept the more likely it may be that this concept is relevant to the active concept.

[00141] In some embodiments, the statistical coherence score may be computed as a function of the ratio of the number of documents containing the active concept, at least one concept in a neighborhood of the active concept, the concept to be scored, and at least one concept in the neighborhood of the concept to be scored (in the numerator) with the number of documents containing the active concept and at least one concept in a neighborhood of the active concept (in the denominator). Calculating a statistical coherence score in this way may be advantageous in that content that includes the label of the concept to be scored, but is directed towards a distinct meaning, is not considered when calculating a statistical coherence score.

[00142] For example, the concepts "field game" and "sport" may be neighbors of the candidate "cricket." Restricting the documents used in computing the statistical coherence score to only the documents that include the concept "cricket" and at least one concept from among "field game" and "sport" may increase the likelihood that documents that include the concept "cricket," but are related to the insect, are avoided when calculating the statistical coherence score.

[00143] In some embodiments, the statistical coherence score may be computed by using only a subset of the documents containing the active concept, at least one concept in the neighborhood of the active concept, the concept to be scored, and, optionally at least one concept in the neighborhood of the concept to be scored. In this case, the statistical coherence score may

be computed as a function of a so-called term frequency (TF) score of the concept to be scored in one or more documents in the aforementioned subset of documents.

[00144] In some embodiments, a TF score for a concept to be scored may be calculated for each document in the subset and the statistical coherence score may be calculated as the average or median of the computed TF scores. Alternatively, the statistical coherence score may be calculated as the largest calculated TF score. This may be advantageous in situations when a concept to be scored appears infrequently in a large number of documents within the subset of documents used for calculating the statistical coherence score.

[00145] In yet another embodiment, the statistical coherence score may be computed as a function of the inverse document frequency (IDF) score, which may be computed as a reciprocal of how frequently a concept to be scored appears within the set of documents used for calculating the statistical coherence score. In yet another embodiment, the statistical coherence score may depend on the product of the term frequency and the inverse document frequency scores. It should be appreciated that, in some embodiments, values calculated in the process of computing the statistical coherence score may be normalized.

[00146] It should be appreciated that, just as the case may be when computing the Jaccard score, the way in which the statistical coherence score is calculated may depend on whether the concept to be scored was retrieved from the semantic network or, instead, was synthesized during act 106 of process 100. In the case that the concept to be scored was synthesized by using an addition operation, the statistical coherence score may be obtained by applying the above-described techniques not to the synthesized concept, but rather to the concept that was combined with the active concept to produce the synthesized concept (e.g., the concept "shiatsu" shown in FIG. 5E). Similarly, when a statistical coherence score is being computed for a concept synthesized via a substitution operation, the statistical coherence score may be obtained by applying the above-described techniques not to the synthesized concept, but to the concept that was used to substitute an attribute of the active concept as part of the substitution (e.g., the concept "press sets" shown in FIG. 5H).

IV.E Cosine Similarity Score

[00147] In the cosine similarity technique, the cosine similarity score of a particular concept may be calculated by using the cosine similarity metric for evaluating semantic proximity between pairs of concepts. To evaluate the cosine similarity metric between two concepts A and B, each of the concepts is mapped to two vectors in Euclidean space of any suitable dimension. The cosine similarity between the two concepts may then be computed as the ratio between the inner product between the two vectors and the product of the magnitudes of the two vectors. This ratio represents the cosine of the angle between the two vectors, giving rise to the name "cosine similarity."

[00148] A concept may be mapped to a vector in any suitable way. For example, a concept may be mapped to a vector comprising a coordinate for each of the concept's attributes, with each coordinate containing a number associated with the attribute. Thus, if concept A has ten attributes, the concept may be mapped to a ten-dimensional vector such that the number in each dimension is associated to the corresponding attribute. The number corresponding to an attribute may be any suitable number and, for example, may be a term frequency (TF) score or a TF-IDF score associated with the attribute.

IV.F Integrated Score

[00149] As previously mentioned, any of the types of scores described above may be combined to form an integrated score that may be assigned to the concepts obtained in act 106 of process 100. Though, in some embodiments, the scores need not be combined and only one of the aforementioned types of scores may be assigned to each concept obtained in act 106.

[00150] In embodiments where one or more types of scores may be combined to form an integrated score, the scores may be combined in any of numerous ways. For example, the scores may be combined by computing a weighted linear combination of the scores to compute the integrated score. The weights used to combine the scores may be any suitable weights and may be increased or decreased to reflect which scores should be weighted more when combining the scores into an integrated score. The scores and/or weights may be normalized in any suitable way prior to be combined into an integrated score.

[00151] After a score is assigned for the concepts obtained in act 106 of process 100, one or more of the scored concepts may be selected for subsequent use based on the calculated scores. The score-based selection may be done in any suitable way. In some embodiments, for example, concepts associated with a score above a predetermined threshold may be selected for subsequent use. Alternatively, a predetermined number or percentage of the top-scoring concepts may be selected. Though, it should be recognized that many other ways of utilizing scores to select one or more concepts will be apparent to one skilled in the art.

V. Provide Content to User(s) Based on Active Concept and Selected Concept(s)

[00152] After one or more concepts relevant to the active concept are selected in act 108, process 100 proceeds to act 110, where information may be provided to the user(s) associated with the user context information obtained in act 102, based at least in part on the active concept identified or generated in act 104 and the relevant concept(s) selected in act 108. To this end, information to present to the user(s) may be selected from among a larger set of information by using the active concept and relevant concept(s) selected in act 108. Though, it should be recognized that the type of information provided to the user(s) and the manner in which the information may be provided to the user(s) may vary depending on the specific scenario in which the techniques described herein may be applied.

[00153] As previously mentioned, in some embodiments, user context information obtained in act 102 may comprise information provided by a user that may indicate the type of information that the user may be interested in. For example, the user context information may comprise a user request for information that the user may be seeking. Such a request may be in any suitable form such as a search query or one or more settings indicating that the user wishes to receive news updates related to a certain topic or topics, advertisements relating to one or more types of product(s), information about updates on any of numerous types of websites, newsletters, e-mail digests, etc. Accordingly, in response to the request, the user may be presented with information obtained, from among a large set of content that may be presented to the user, based at least in part on the active concept, which was derived from the user's request, and the relevant concepts to the active concept that were selected in act 108.

[00154] For example, if the user's request comprised a search query, the active concept and the related concepts, selected in act 108, may be used to generate one or more search queries to be provided to one or more search services. This may be done in any suitable way. For example, a search query may be constructed from the active concept and any of the selected concepts by using the labels and attributes associated with these concepts. A search query may be formed by joining the concept labels and attributes from the active concepts and any of the selected concepts by using various Boolean operators such as "AND" and "OR." For example, if the active concept representing a user's search query is the concept "yoga mat," described with reference to FIG. 5F and the concept "sweat-absorbent yoga mat" is selected in act 108, a search query "(yoga mat) AND (sweat-absorbent)" may be formed. As another example, if the active concept representing a user's search query is the concept "recliner," described with reference to FIG. 5D and the concept "massage recliner" is selected in act 108, a search query "(recliner) OR (massage recliner)" may be formed. Other more complex search queries may be formed and may include keywords associated with multiple selected concepts, any disambiguation terms used to identify the active concept, and/or any other user context information. For example, suppose a user is searching for an Italian restaurant at 10pm, while renting a car at the airport. The active concept "Italian restaurant" may be used to select relevant concepts such as "Pizza," "Pasta," and "Carbs" and, together with geo-spatial information about the user obtained from the user context information, be used to construct a query such as "(Italian Restaurant) OR (Pizza) OR (Pasta) OR (Carbs) AND (New York) AND (Airport) and (OPEN AFTER 10pm)."

[00155] These types of complex queries would rarely be composed by users. Such queries create an effective semantic search, even if the content has not been semantically analyzed in advance (e.g., unstructured content), because such a query will match literal terms in the content indexed by the search service that are not necessarily literal terms in the original query.

[00156] The search service may be any general-purpose search engine. For instance, the search service may be any search engine that may be publicly accessible via the Internet. As another example, the search service may be a search engine accessible via any computer network other than the Internet. Examples of such search engines include search engines used for searching a corporate intranet or any other private network.

[00157] In response to issuing the one or more search queries to the search service, a set of search results may be received from the search service. The text (or fragments of the text) of the documents or pieces of content in the search results may be compared to the active concept and/or the concept(s) selected in act 108 and the returned search results may be ranked and/or filtered out based on how closely they match these concept definitions.

[00158] Any of a variety of possible ranking or filtering techniques may be used, as the invention is not limited in this respect. However, such techniques may enable the provisioning of content to users without overwhelming the users with information irrelevant to the users. Search services may provide a number of textual features in their search results: titles, abstracts, descriptions, tags, hyperlinks, etc. These textual features may provide for text analysis as a means to filter the search engine results against the terms provided through concepts selected in act 108, for example, by comparing the terms against words in the textual features of the search engine results. Whole or partial matches of terms may be used to weight the relevance of the individual results. In some embodiments, the search results returned from the search service may not include the identified pieces of content themselves, but rather may include a list of hyperlinks to these pieces of content along with an excerpt of each piece of content. In such embodiments, rather than retrieving each piece of content using the provided hyperlink, the list of hyperlinks may be filtered and ranked using the associated excerpt, and the excerpt may be semantically annotated.

[00159] In some embodiments, user context information obtained in act 102 may comprise information related to a user that may indicate the type of information that the user may be interested in. For example, information related to the user may comprise demographic information, the user's Internet browsing history, any information associated with the user on a website such as a social networking website, geo-spatial information may comprise the current location of the user's computing, etc. Accordingly, the user may be presented with information obtained at least in part based on the active concept, which was derived from information related to the user, and the selected concepts.

[00160] For example, a user may be presented with personalized product and service recommendations based on the active concept and the selected concepts. Consequently, the

personalized recommendations may reflect one or more of the user's interests. The personalized recommendations may include promotional content including, but not limited to, advertisements for products and/or services. For example, an active concept, derived from user context information, may indicate that the user is interested in "recliners." Accordingly, the user may be presented with advertisements related to "recliners" and to "massage recliners," which is a concept relevant to the active concept "recliners" as described with reference to FIG. 5D. As such, the user may not be presented with irrelevant promotional content.

[00161] As another example, a user may use an online information source (or multiple websites) to obtain information that the user may be interested in. The online information source may be any suitable information source and, for example, may be an Internet portal, an Intranet portal, a news website, a social networking website, a micro-blogging service, a blog service, a blog reader, a shopping website, real-time feeds, etc. Each such online information source may be configured to present any of numerous types of information to a user including, but not limited to, news, advertisements, content recommendations, real-time updates (e.g., tweets). As such, when the user uses the online information source, the user may not be overwhelmed with irrelevant content.

[00162] Accordingly, in some embodiments, the active concept and the selected concepts may be used to rank, prioritize and/or filter the information that may be presented to a user such that the information presented to the user may reflect that user's interests. This may be done in any suitable way. For example, any of the information that a website may be configured to present to a user may comprise one or more textual features (e.g., tags, text, hyperlinks, descriptions, etc.). These textual features may be compared to any keywords provided through the active concept and concepts selected in act 108. Whole or partial matches may be used to weight the relevant of the individual terms.

VI. Performing a Semantic Operation on a Digital Social Network

[00163] Online social networking has seen unprecedented growth over the last number of years. Recent technologies have enabled establishing connections between thousands, millions,

or hundreds of millions of people around the world over a private and public networks in ways previously unimagined. With over a billion registered users participating online at hundreds of social networking sites that exist today, online social networking sites have grown tremendously in importance as a venue for finding and sharing information with others.

[00164] However, finding relevant information on a digital social network that may be of interest to a particular user presents many challenges. Literal matches between user-performed search queries for groups, discussions, status updates, tweets, relevant connections, other social network users, jobs, blogs, news, events, etc. often presents an overwhelming amount of information that the user must manually sift through. Prioritizing amongst a large number of results, thus becomes important so that the user may experience more immediately experience the most relevant sought-after material. For example, a search for professionals that have experience in "computer developer" may produce in the search results an unmanageable number of members from the social network with such experience. Should the user performing this search be more interested developers with experience in a certain computer programming language, such as Java developers rather than C+ developers, or be more interested in developers within a given industry, such as risk-analysis rather than mobile applications, such preferences could be used to prioritize the voluminous search results.

[00165] Furthermore, there may be a disparity between the terminology employed by a user and the terminology found in items that are of interest to the user. This gap may be a particularly felt when recommendations are based on descriptions that the user has provided in his or her profile. For example, though a user may have listed "applied mathematics research" in the skills section of their professional network profile, performing a literal match search for groups would omit a group directed towards "quantitative modeling," despite such a group having a reasonable likelihood of being of interest to that user. Accordingly, information retrieval problems exist on both the end of an inability to find sufficient relevant information that as well as the inundation of too much information that is only marginally relevant to the user's interests.

[00166] The inventors of the present invention have realized that the aspects of techniques and methodologies explained above may be leveraged to address these and related

problems that may be experienced in the realm of digital social networks. As will now be described, various embodiments of the present system and method may be utilized to perform a semantic operation on a social network to filter, rank or augment a retrieval of information relevant to a user of the social network. More generally, in accordance with one embodiment, a system and method comprises first receiving a social network user context information associated with a user of the social network. Here, a use of a social network refers to any individual that visits or avails themselves to the benefits of a social network. The user may be a registered member, or may be a non-registrant of the social network who may still browse information on the social network anonymously, or under a temporary "guest" membership.

[00167] In an embodiment, a user's social connections may be determined by looking at a list of connections that a user has within a given social network, and if possible, also looking at a list of connections that a user may have within one or more other social networks. If a user has multiple sets of connections across a number of different social networks, then potentially all of those connections and any cross-connections identified between them may be utilized as an input into development of the user context. If the user's various social connections on one or more social networks have some features or characteristics in common, then those common features or characteristics may also be used as an input into developing the user context.

[00168] In an embodiment, discussions in which a user participates may be utilized as an input into development of a user context, and this information may also be used to bring in supplemental information related to the discussion. For example, if a user participates in a discussion on a recent political event in particular country, then news articles relating to that event may be searched for and pulled in to create a better picture of what the discussion relates to.

[00169] Topics of interest or trends as identified by a user's connections within one or more social networks may also form a part of the user context. For example, a user may be influenced by the developing trends or opinions as expressed by a majority of people within their social network, as they are opinions expressed by people the user personally knows and may implicitly trust. If a majority of people within a user's social network "like," "dig," or status update post about a topic, for example a particular tire brand for winter driving, this may result in

the tire brand becoming more desirable to a user based on one or more recommendations from a trusted contact.

[00170] Various e-commerce activities undertaken by a user may also form another input into a user context. For example, if a user has a number of favourite fashion or clothing websites at which the user regularly shops, such fashion and clothing websites may be noted and form a part of the user's detailed social graph as preferences for particular fashion brands or for particular online retail establishments.

[00171] Various multimedia content may also be utilized as inputs into a user's social context. The multimedia content may comprise videos, images, and audio. Images and videos posted on a user's home webpage may be utilized as inputs, for example. Any video input can then also be processed by associated content/meta tags and/or speech, images or text within the video. Audio input from multimedia files posted on a user's social page may also provide input into a user's social graph. Meta tags associated with the multimedia content may also be used to identify a user's preferences. For example, if the audio inputs are repeatedly from a certain artist, then the artist's name may be utilized to form a part of the user's social graph for targeted advertising and promotions.

[00172] The various types of inputs described above may be synthesized in order to obtain a more accurate user context. This may include identifying multiple interests, and also determining what preferences a user has in a social networking context by synthesizing inputs obtained from the user's participation in one or more social networks.

[00173] Based on the profile of a user obtained from synthesizing various inputs based on a user's online social networking interactions, the system and method of the present disclosure may present information better matching the user's interests.

[00174] Now referring to FIG. 7, in accordance with an embodiment, shown is an illustrative process 700 for performing a semantic operation on a digital social network. The process of FIG. 7 begins at act 702, where a user context information may be obtained from a social network. The user context information may comprise information about the user that may be used to ascertain the information that the user may be interested in.

[00175] The types of inputs from a user's online interactions that may be used to obtain a user context may include, for example: (i) a user's social connections in one or more social networks, (ii) discussions that a user is participating in on various chat rooms or blogs on the social network, (iii) topics of interest as indicated by selections or postings on a user's social network profile, (iv) a user's e-commerce activities, and (v) multimedia content viewed by or uploaded to a user's social webpage or blog on the social network. With these various types of inputs, all of the techniques for obtaining context information, as described above in section "I. Obtaining User Context Information" and act 102 of process 100, may be applied to obtain a user context.

[00176] Still referring to FIG. 7, following act 702, process 700 next proceeds to act 704, where process 700 identifies or generates an active concept representing the user context information in a Knowledge Representation. Act 704 in many aspects parallels act 104 of process 100 and may employ any of the techniques described above in the section "II. Identifying or Generating Active Concept Representing User Context Information".

[00177] Still referring to FIG. 7, act 706 illustrates obtaining one or more concepts relevant to an active concept based on a KR. Act 708 of process 700 entails calculating a value based on the score of the selected concepts. The scoring and selecting of obtained relevant concepts may essentially parallel the methodologies illustrated in acts 106 and 108 of Process 100 respectively, and described in greater detail above in the section "III. Identifying Concepts Relevant to Active Concept" and "IV. Score and Select Identified Concept(s) by using Relevance Measure or Measures" respectively. The KR used for obtaining semantically relevant concepts in item 706 may be constructed based on any content, though in some embodiments the content used to form the KR may be directed towards one or more particular subjects pertaining to a user's current online interaction.

[00178] After the relevant concepts have been scored and the top scoring concepts have been selected, process 700 then proceeds to act 710, where process 700 uses the top scoring concepts to establish or add to an interest network with the top scoring concepts. With the establishment of such an interest network, process 700 can then proceed to act 712, where

process 700 filters, ranks (or re-ranks), or augments retrieval of information based on the interest network.

[00179] FIG. 8 illustrates one non-limiting example where process 700 may be carried out to establish an interest network for a user named Johnny Appleseed. In his user profile, Johnny Appleseed has indicated that his specialities are "corporate strategy", "quantitative modelling", and the "automotive industry". Also listed on Johnny Appleseed's profile are a number of hobbies in which he is interested.

[00180] The specialities, hobbies and other information recited on the user-profile shown in FIG. 8 provide a social-network user-context, which is exemplary of the information obtained in act 702 of FIG. 7. FIG. 8 illustrates how, in one aspect, this user-context information allows for identifying or generating the active concept "corporate strategy" in a KR, which demonstrates act 706 of process 700. The active concept may have already existed in the KR or alternatively may have been generated as described with respect to FIG. 3 and its associated description. The active concept may then be used to obtain a plurality of relevant concepts in the KR as recited by act 706 of act 700, which is shown in FIG. 8 with the active concept "corporate strategy" being mapped to concepts in the KR such as "military strategy", "business", etc. From the plurality of relevant concepts in the KR, a number of concepts have been selected based on their relative high scores, depicting acts 708 and 710 of process 700. In this illustrative example, the phrases "business principles", "corporate values" and "management strategy" have all been selected as top-scoring concepts.

[00181] The selected concepts may then used to establish or add to an interest graph. As shown in FIG. 8, the interest graph includes "corporate strategy" as a concept (taken as an explicit recitation from the user's profile) and has newly formed relationships to concepts including "management strategy", "business principles", and "corporate values". Thus, in some aspects, the interest graph may include the active concept connected to as newly retrieved or synthesized concepts obtained from that active concept.

[00182] Further, "corporate strategy" and its associated connections may be joined to a virtual node, which may be created to link portions of the interest graph with one another. In the example, "automotive industry" was another specialization explicitly recited by the user.

"Automotive industry" was used to previously create, using process 700, an interest graph with connections to "lean manufacturing", "car design", and "fuel economy". By forming a virtual node, in this case labeled with the name of the user ("Johnny Appleseed") in FIG. 8, two different interest graphs based on different explicit recitations by the user may be adjoined.

[00183] At this point, the concepts in the interest graph may be used to rank, filter or augment information in the social network that may be relevant to the user. FIG. 8 shows an ordering of relevant discussion groups: "1. Applying Corporate Strategy to Lean Manufacturing", "2. Management Strategy in the Automobile Industry", "3. Corporate Strategy for Beginners", etc. It should be noted that while items "1." and "3." each recite "corporate strategy", item "1." may be prioritized higher due to it's additional recitation of the "lean manufacturing" as obtained from the user's interest graph. Similarly, results may be filtered out such that only the results that recite the most concepts in the user's interest graph may appear. Because the interest graph also adds concepts that previously were not known to be relevant to the user (e.g. "business principles"), the interest graph may be employed to augment results that would otherwise be sparse or null. For example, if no items in the social network recited "corporate strategy" but some items "business principle," those items that recited "business principles" may be presented to the user potentially relevant.

[00184] While the example in FIG. 8 illustrates suggested groups for the user, the interest graph may be employed to rank, filter or augment any type of information that may be relevant to the user of the digital social network. For example, the interest graph may be used to rank, filter or augment information such as news, suggested connections to other social network users, status updates, prospective jobs listings, etc. Thus, with the application of the present system and method, it will be seen that based on a user context obtained for social networking context of a user that information experienced by the social network user may be enhanced.

[00185] In an embodiment of the present system and method, a user's account such as found at an online retailer may be linked directly to a user context in a dynamic manner, such that a given user context is always up-to-date in terms of what the user is most currently interested in. As noted earlier, the inputs to a user context may be of many different types, and come from a number of different sources including multiple social networking user accounts.

This diversified source of information may provide a more accurate, dynamic, and balanced profile of the user at any given moment.

VII. Additional Implementation Detail

[00186] The above-discussed computing devices (e.g., client computer and server shown in FIGS. 2A and 2B) may be implemented in any of a variety of ways. FIG. 9 is a block diagram an illustrative computing device 1000 that may be used to implement any of the above-discussed computing devices.

[00187] The computing device 1000 may include one or more processors (e.g., microprocessors) 1001 and one or more tangible, non-transitory computer-readable storage media (e.g., memory 1003). Memory 1003 may store, in tangible non-transitory computer-readable storage media computer instructions that implement any of the above-described functionality. Processor(s) 1001 may be coupled to memory 1003 and may execute such computer instructions to cause the functionality to be realized and performed. Computing device 1000 may also include a network input/output (I/O) interface 1005 via which the computing device may communicate with other computers (e.g., over a network). In some embodiments, the computing device may also include one or more user I/O interfaces, via which the computer may provide output to and receive input from a user. The user I/O interfaces may include devices such as a keyboard, a mouse, a microphone, a display device (e.g., a monitor or touch screen), speakers, a camera, and/or various other types of I/O devices.

[00188] The above-described embodiments of the present invention can be implemented in any of numerous ways. For example, the embodiments may be implemented using hardware, software or a combination thereof. When implemented in software, the software code may be embodied as stored program instructions that may be executed on any suitable processor or collection of processors (e.g., a microprocessor or microprocessors), whether provided in a single computer or distributed among multiple computers.

[00189] It should be appreciated that a computer may be embodied in any of numerous forms, such as a rack-mounted computer, a desktop computer, a laptop computer, or a tablet computer. Additionally, a computer may be embodied in a device not generally regarded as a

computer, but with suitable processing capabilities, including a Personal Digital Assistant (PDA), a smart phone, a tablet, a reader, or any other suitable portable or fixed electronic device.

[00190] Also, a computer may have one or more input and output devices. These devices may be used, among other things, to present a user interface. Examples of output devices that may be used to provide a user interface include printers or display screens for visual presentation of output, and speakers or other sound generating devices for audible presentation of output. Examples of input devices that may be used for a user interface include keyboards, microphones, and pointing devices, such as mice, touch pads, and digitizing tablets.

[00191] Such computers may be interconnected by one or more networks in any suitable form, including networks such as a local area network (LAN) or a wide area network (WAN), such as an enterprise network, an intelligent network (IN) or the Internet. Such networks may be based on any suitable technology and may operate according to any suitable protocol and may include wireless networks, wired networks, and/or fiber optic networks.

[00192] The various methods or processes outlined herein may be coded as software that is executable on one or more processors that employ any one of a variety of operating systems or platforms. Additionally, such software may be written using any of numerous suitable programming languages and/or programming or scripting tools, and also may be compiled as executable machine language code or intermediate code that is executed on a virtual machine or a suitable framework.

[00193] In this respect, various inventive concepts may be embodied as at least one non-transitory tangible computer-readable storage medium (e.g., a computer memory, one or more floppy discs, compact discs, optical discs, magnetic tapes, flash memories, circuit configurations in Field Programmable Gate Arrays or other semiconductor devices, etc.) article(s) encoded with one or more programs that, when executed on one or more computers or other processors, implement the various process embodiments of the present invention. The non-transitory computer-readable medium or media may be transportable, such that the program or programs stored thereon may be loaded onto any suitable computer resource to implement various aspects of the present invention as discussed above.

[00194] The terms "program" or "software" are used herein in a generic sense to refer to any type of computer code or set of computer-executable instructions that can be employed to program a computer or other processor to implement various aspects of embodiments as discussed above. Additionally, it should be appreciated that according to one aspect, one or more computer programs that when executed perform methods of the present invention need not reside on a single computer or processor, but may be distributed in a modular fashion among different computers or processors to implement various aspects of the present invention.

[00195] Computer-executable instructions may be in many forms, such as program modules, executed by one or more computers or other devices. Generally, program modules include routines, programs, items, components, data structures, etc. that perform particular tasks or implement particular abstract data types. Typically the functionality of the program modules may be combined or distributed as desired in various embodiments.

[00196] Also, data structures may be stored in non-transitory tangible computer-readable storage media articles in any suitable form. For simplicity of illustration, data structures may be shown to have fields that are related through location in the data structure. Such relationships may likewise be achieved by assigning storage for the fields with locations in a non-transitory tangible computer-readable medium that convey relationship between the fields. However, any suitable mechanism may be used to establish relationships among information in fields of a data structure, including through the use of pointers, tags or other mechanisms that establish relationships among data elements.

[00197] Also, various inventive concepts may be embodied as one or more methods, of which multiple examples have been provided (e.g., processes 100, 300, and 700). The acts performed as part of the method may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments, or vice versa.

[00198] All definitions, as defined and used herein, should be understood to control over dictionary definitions, definitions in documents incorporated by reference, and/or ordinary meanings of the defined terms.

[00199] The indefinite articles "a" and "an," as used herein, unless clearly indicated to the contrary, should be understood to mean "at least one."

[00200] As used herein, the phrase "at least one," in reference to a list of one or more elements, should be understood to mean at least one element selected from any one or more of the elements in the list of elements, but not necessarily including at least one of each and every element specifically listed within the list of elements and not excluding any combinations of elements in the list of elements. This definition also allows that elements may optionally be present other than the elements specifically identified within the list of elements to which the phrase "at least one" refers, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, "at least one of A and B" (or, equivalently, "at least one of A or B," or, equivalently "at least one of A and/or B") can refer, in one embodiment, to at least one, optionally including more than one, A, with no B present (and optionally including elements other than B); in another embodiment, to at least one, optionally including more than one, B, with no A present (and optionally including elements other than A); in yet another embodiment, to at least one, optionally including more than one, A, and at least one, optionally including more than one, B (and optionally including other elements); etc.

[00201] The phrase "and/or," as used herein, should be understood to mean "either or both" of the elements so conjoined, i.e., elements that are conjunctively present in some cases and disjunctively present in other cases. Multiple elements listed with "and/or" should be construed in the same fashion, i.e., "one or more" of the elements so conjoined. Other elements may optionally be present other than the elements specifically identified by the "and/or" clause, whether related or unrelated to those elements specifically identified. Thus, as a non-limiting example, a reference to "A and/or B", when used in conjunction with open-ended language such as "comprising" can refer, in one embodiment, to A only (optionally including elements other than B); in another embodiment, to B only (optionally including elements other than A); in yet another embodiment, to both A and B (optionally including other elements); etc.

[00202] As used herein, "or" should be understood to have the same meaning as "and/or" as defined above. For example, when separating items in a list, "or" or "and/or" shall be

interpreted as being inclusive, i.e., the inclusion of at least one, but also including more than one, of a number or list of elements, and, optionally, additional unlisted items.

[00203] The phraseology and terminology used herein is for the purpose of description and should not be regarded as limiting. The use of "including," "comprising," "having," "containing", "involving", and variations thereof, is meant to encompass the items listed thereafter and additional items.

[00204] Having described several embodiments of the invention in detail, various modifications, and improvements will readily occur to those skilled in the art. Such modifications and improvements are intended to be within the spirit and scope of the invention. Accordingly, the foregoing description is by way of example only, and is not intended as limiting.

CLAIMS:

1. A computer-implemented method for performing a semantic operation on a social network, the method comprising:

receiving a social network user context associated with a user of the social network;

generating, through a semantic operation, an interest network based on the user context information; and

filtering, ranking or augmenting, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network;

wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network.

2. The computer-implemented method of claim 1, further comprising representing the interest network as an interest graph.

3. The computer-implemented method of claim 1, wherein the semantic operation is a synthesis operation or retrieval operation performed on a knowledge representation.

4. The computer-implemented method of claim 3, wherein the knowledge representation comprises a semantic network.

5. The computer-implemented method of claim 1, wherein the information is at least one of groups, jobs, discussions, news, products, SMS, applications or other users.

6. The computer-implemented method of claim 1, further comprising scoring the information related to the social network based on an occurrence of at least a portion of the interest network within the information.

7. A system for performing a semantic operation on a social network, the system adapted to:

receive a social network user context associated with a user of the social network;

generate, through a semantic operation, an interest network based on the user context information; and

filter, rank or augment, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network;

wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network.

8. The system of claim 7, wherein the system is further adapted to represent the interest network as an interest graph.

9. The system of claim 7, wherein the semantic operation is a synthesis operation or retrieval operation performed on a knowledge representation.

10. The system of claim 9, wherein the knowledge representation comprises a semantic network.

11. The system of claim 7, wherein the information is at least one of groups, jobs, discussions, news, products, SMS, applications or other users.

12. The system of claim 7, wherein the system is further adapted to score the information related to the social network based on an occurrence of at least a portion of the interest network within the information.

13. A non-transitory computer-readable medium storing computer code that when executed on a computer device adapts the device to perform a semantic operation on a social network, the computer-readable medium comprising:

code for receiving a social network user context associated with a user of the social network;

code for generating, through a semantic operation, an interest network based on the user context information; and

code for filtering, ranking or augmenting, using at least one processor executing stored program instructions, a retrieval of information related to the social network based on the interest network;

wherein the interest network comprises concepts represented by a data structure associated with the concepts in the interest network.

14. The computer-readable medium of claim 13, further comprising code for representing the interest network as an interest graph.

15. The computer-readable medium of claim 13, wherein the semantic operation is a synthesis operation or retrieval operation performed on a knowledge representation.

16. The computer-readable medium of claim 15, wherein the knowledge representation comprises a semantic network.

17. The computer-readable medium of claim 13, wherein the information is at least one of groups, jobs, discussions, news, products, SMS, applications or other users.

18. The computer-readable medium of claim 13, further comprising code for scoring the information related to the social network based on an occurrence of at least a portion of the interest network within the information.

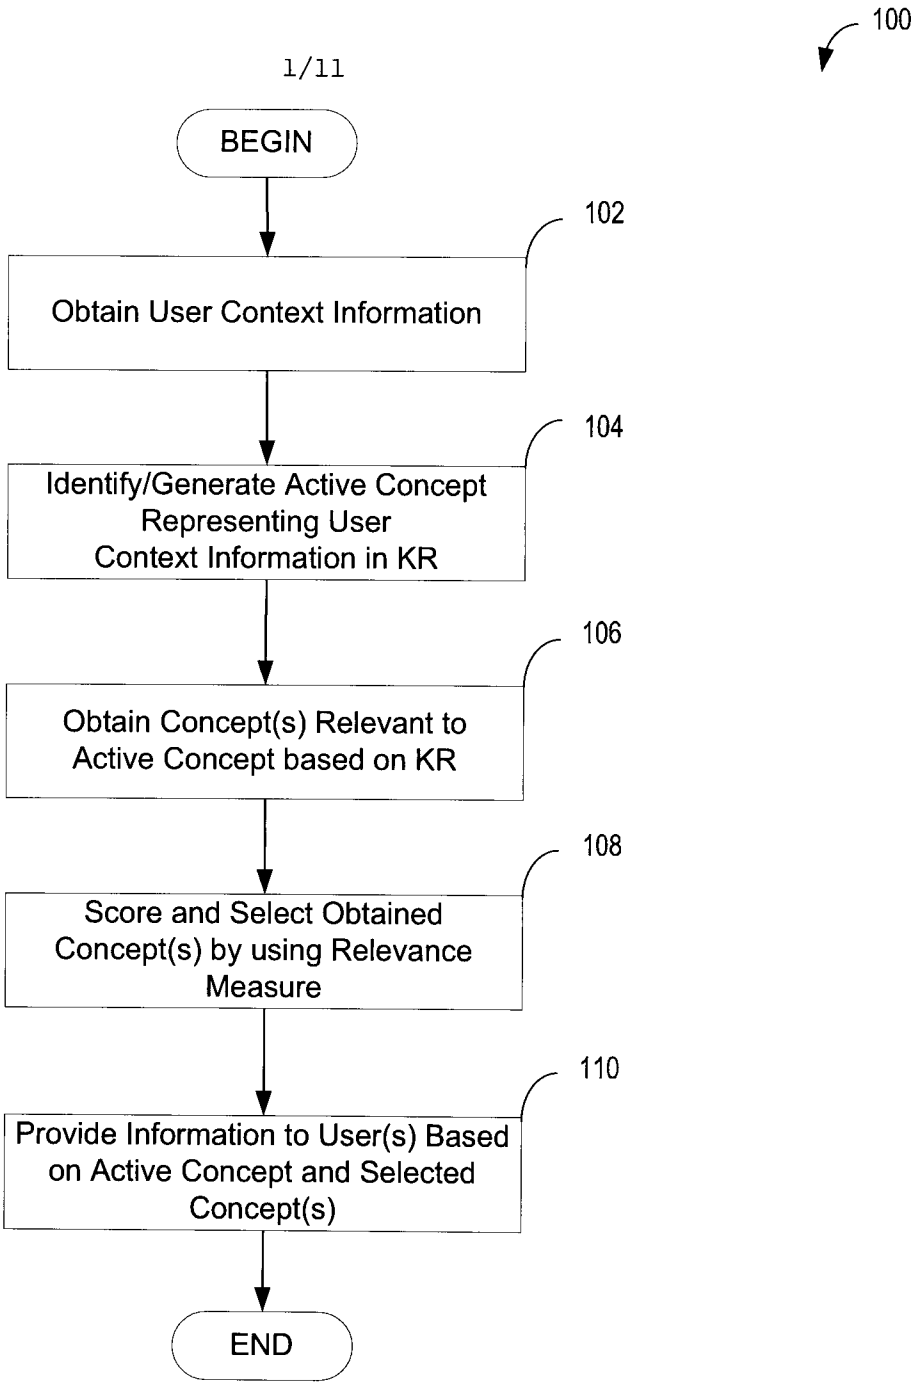
100
BEGIN
102
Obtain User Context Information
104
Identify/Generate Active Concept Representing User Context Information in KR
106
Obtain Concept(s) Relevant to Active Concept based on KR
108
Score and Select Obtained Concept(s) by using Relevance Measure
110
Provide Information to User(s) Based on Active Concept and Selected Concept(s)
END

FIG. 1

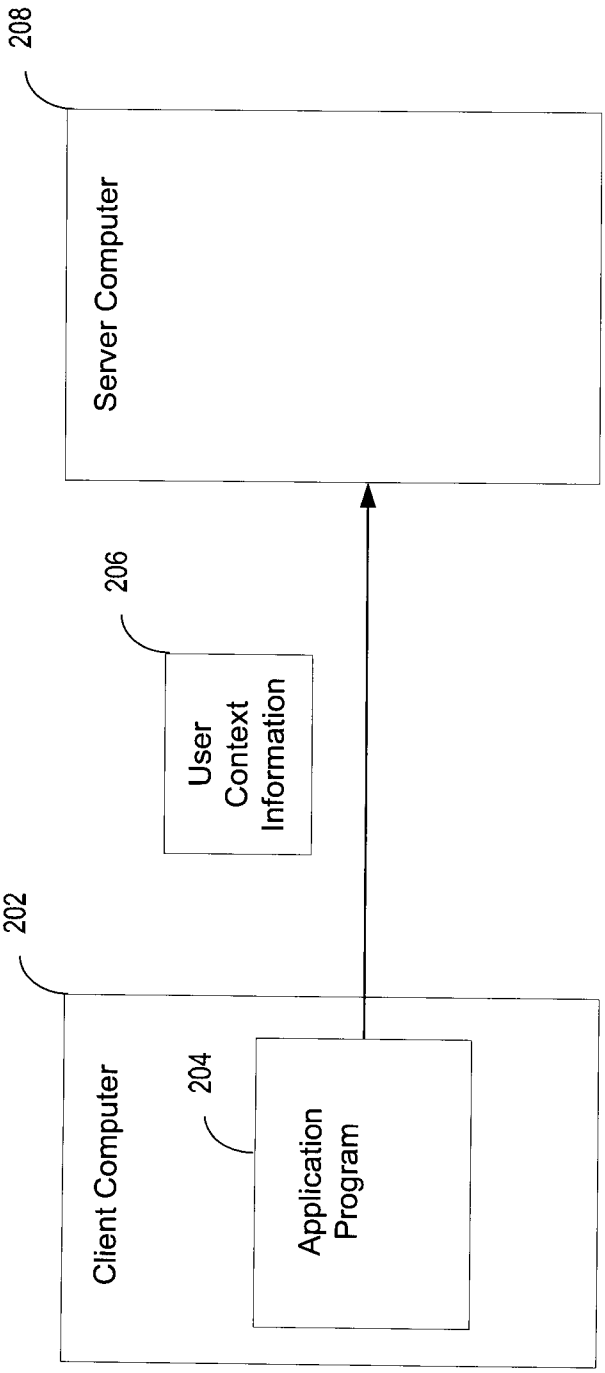
208
Server Computer
206
User Context Information
202
Client Computer
204
Application Program

FIG. 2

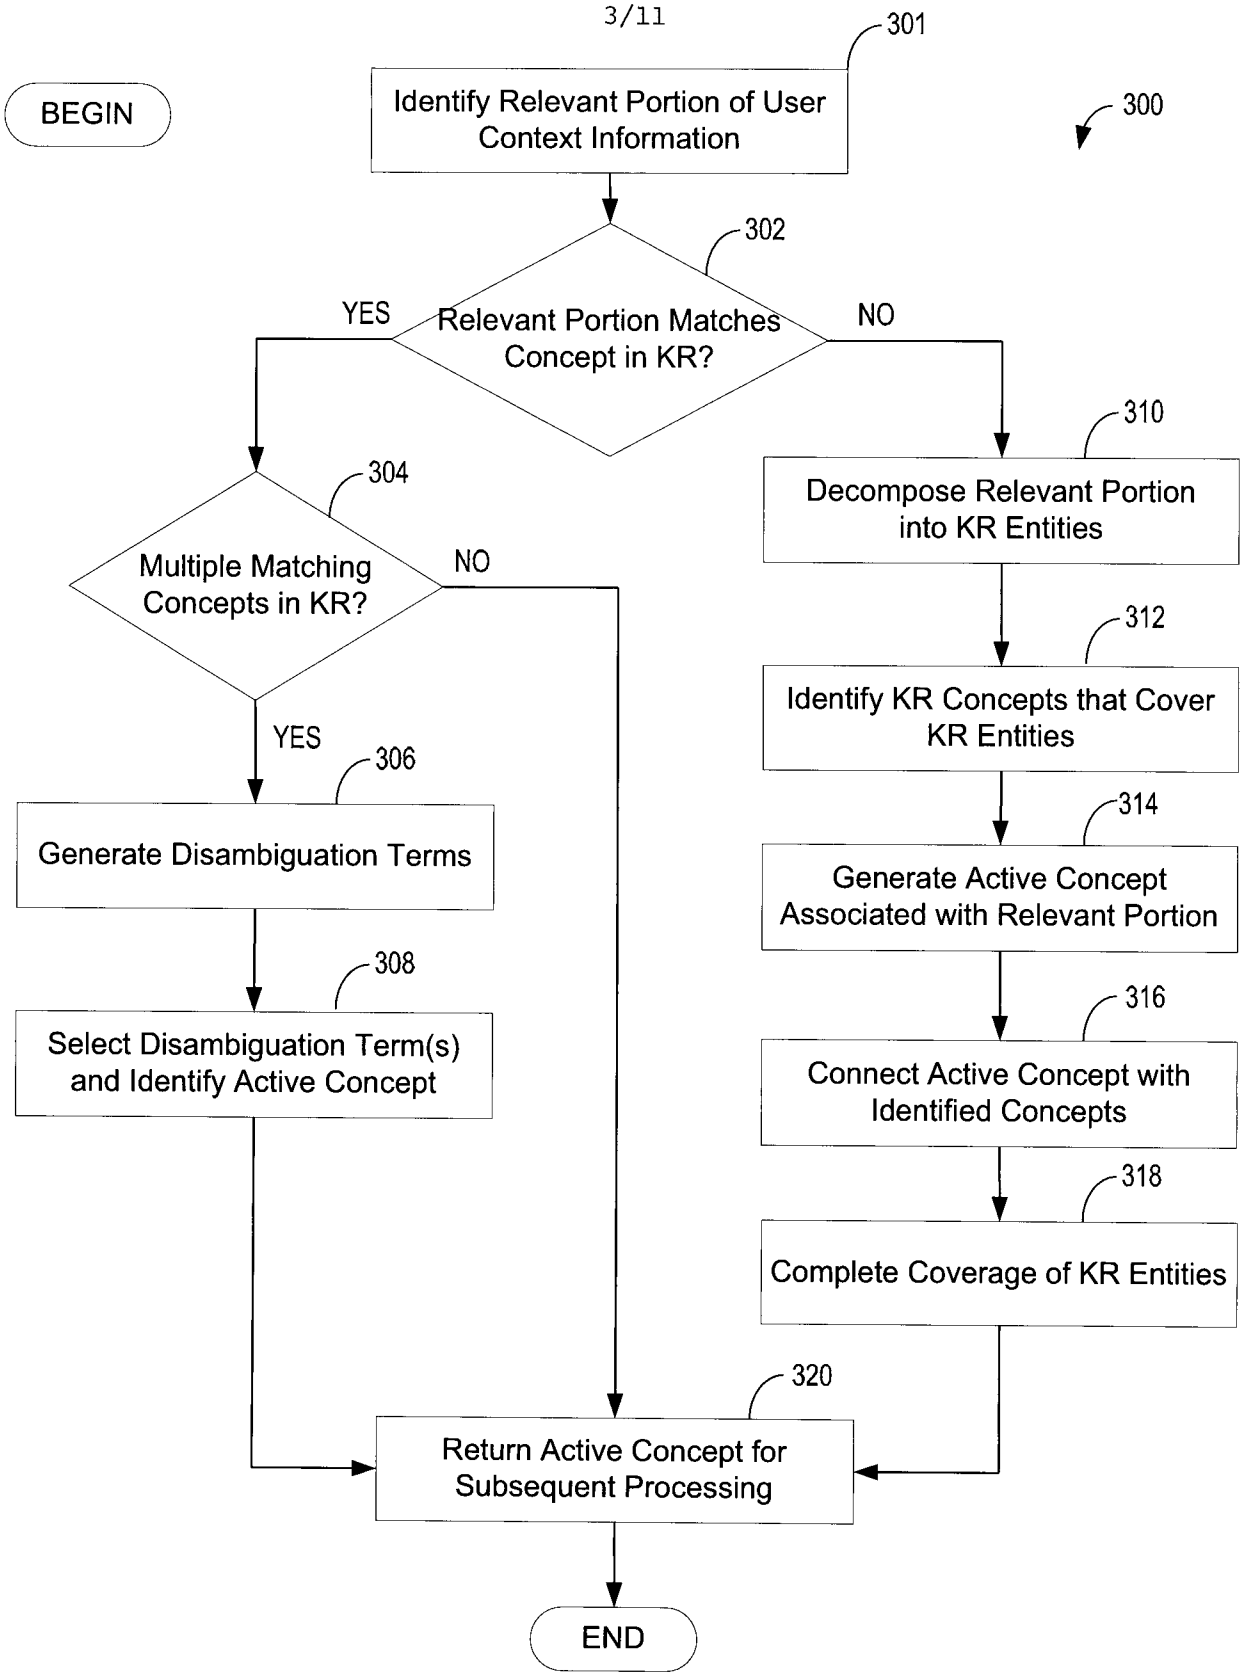
BEGIN
301
Identify Relevant Portion of User Context Information
300
302
Relevant Portion Matches Concept in KR?
YES
NO
304
Multiple Matching Concepts in KR?
NO
YES
306
Generate Disambiguation Terms
308
Select Disambiguation Term(s) and Identify Active Concept
310
Decompose Relevant Portion into KR Entities
312
Identify KR Concepts that Cover KR Entities
314
Generate Active Concept Associated with Relevant Portion
316
Connect Active Concept with Identified Concepts
318
Complete Coverage of KR Entities
320
Return Active Concept for Subsequent Processing
END

FIG. 3

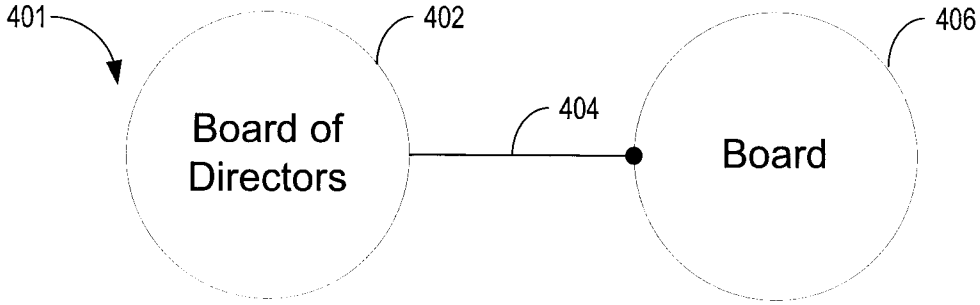

FIG. 4A

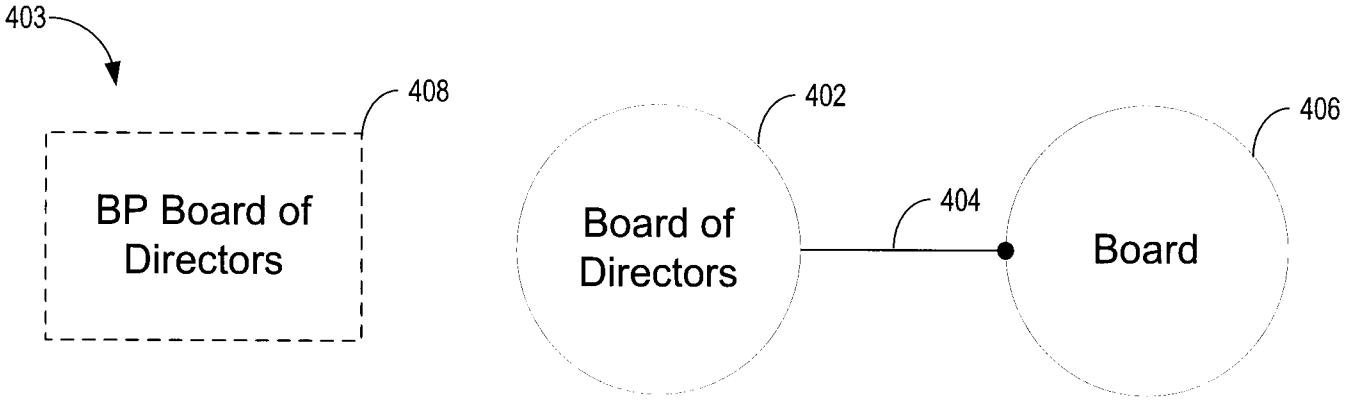

FIG. 4B

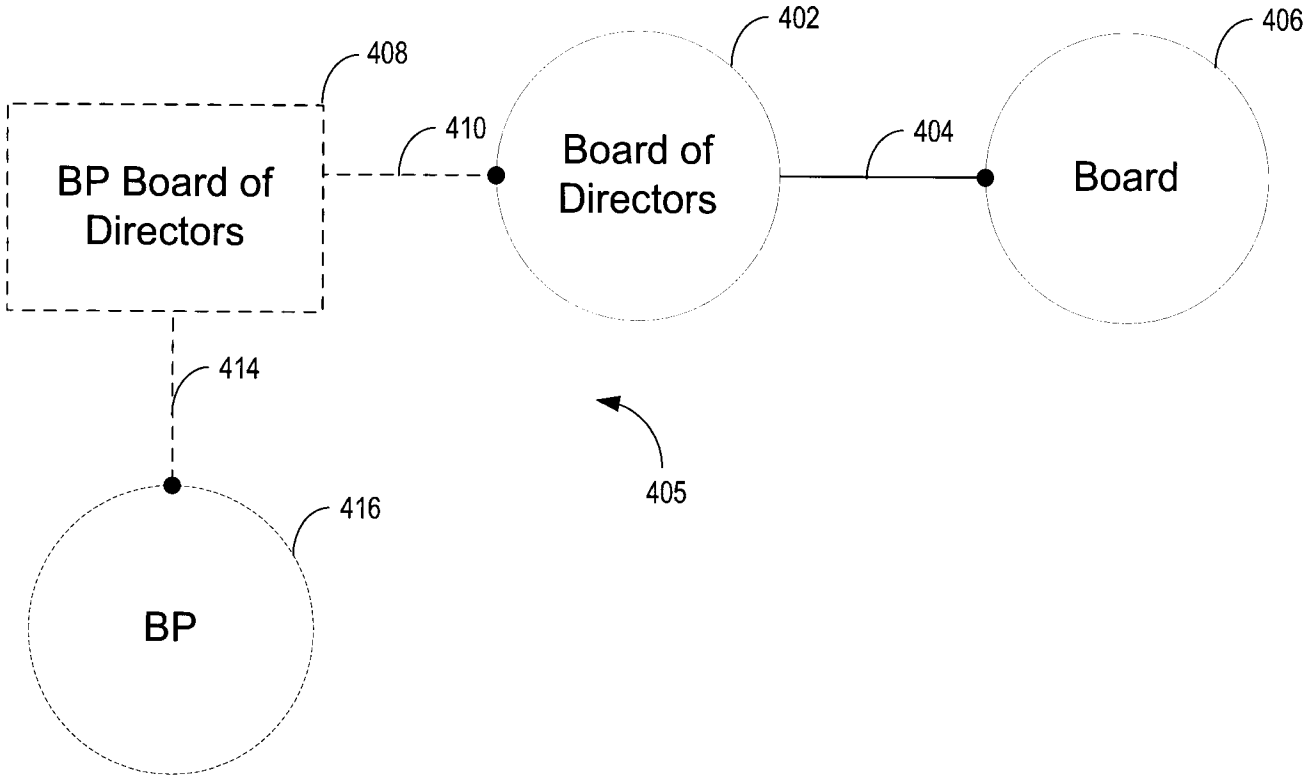

FIG. 4C

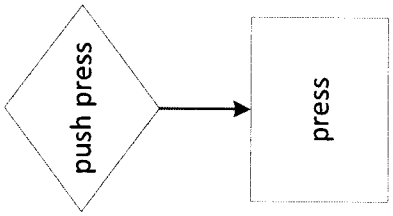

FIG. 5A

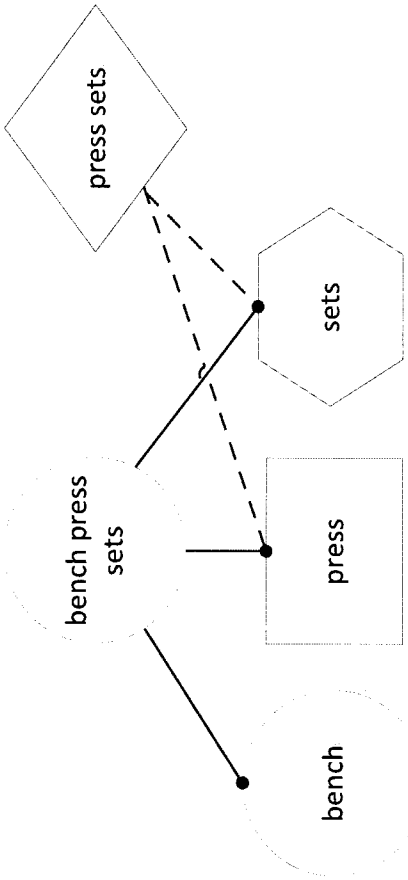

FIG. 5B

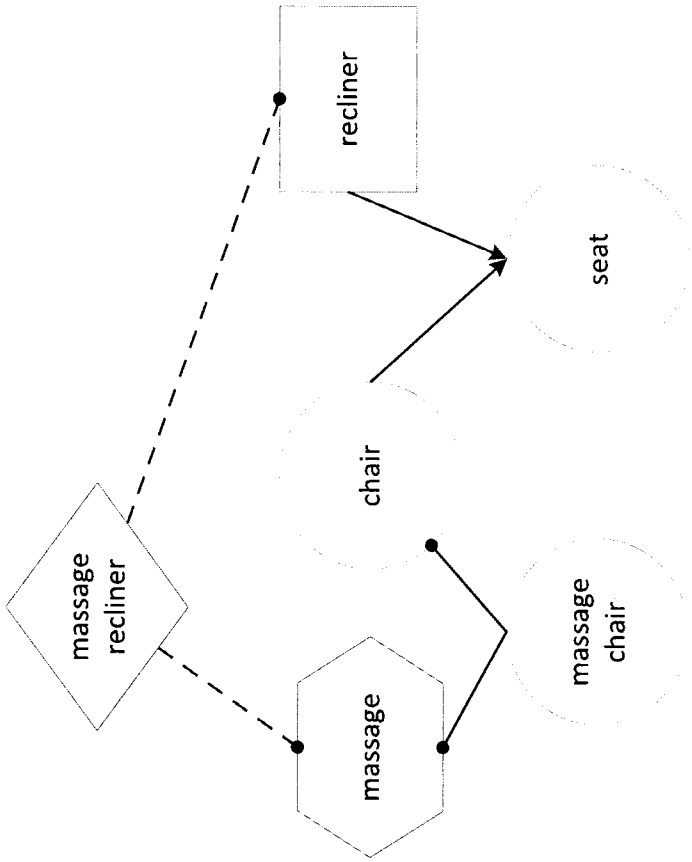

FIG. 5D

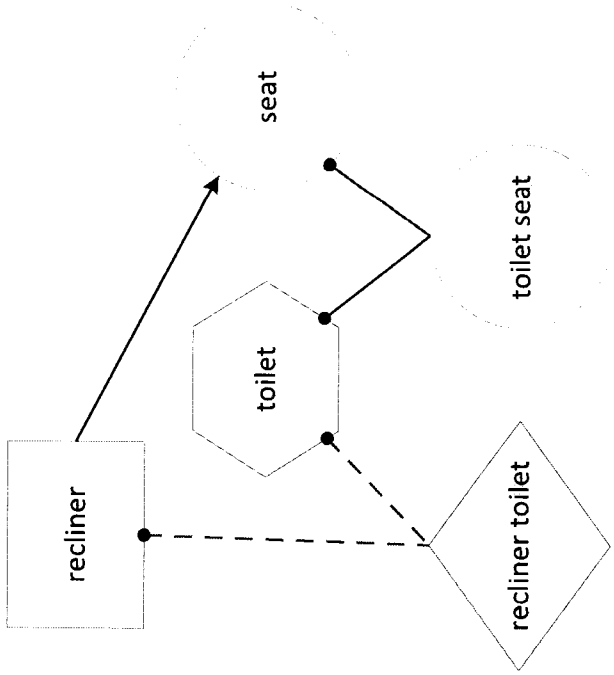

FIG. 5C

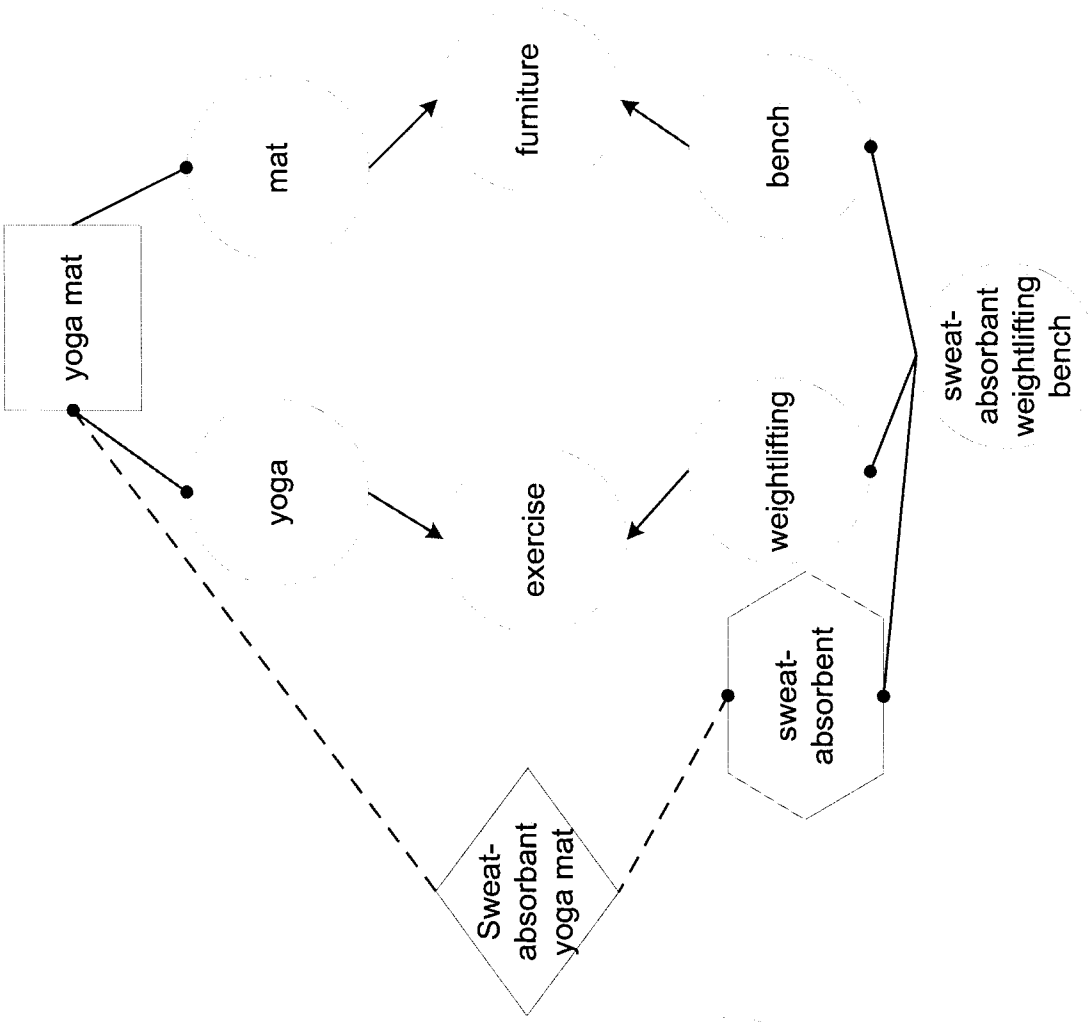
yoga mat
mat
furniture
bench
yoga
exercise
weightlifting
sweat-absorbant weightlifting bench
sweat-absorbent
Sweat-absorbant yoga mat

FIG. 5F

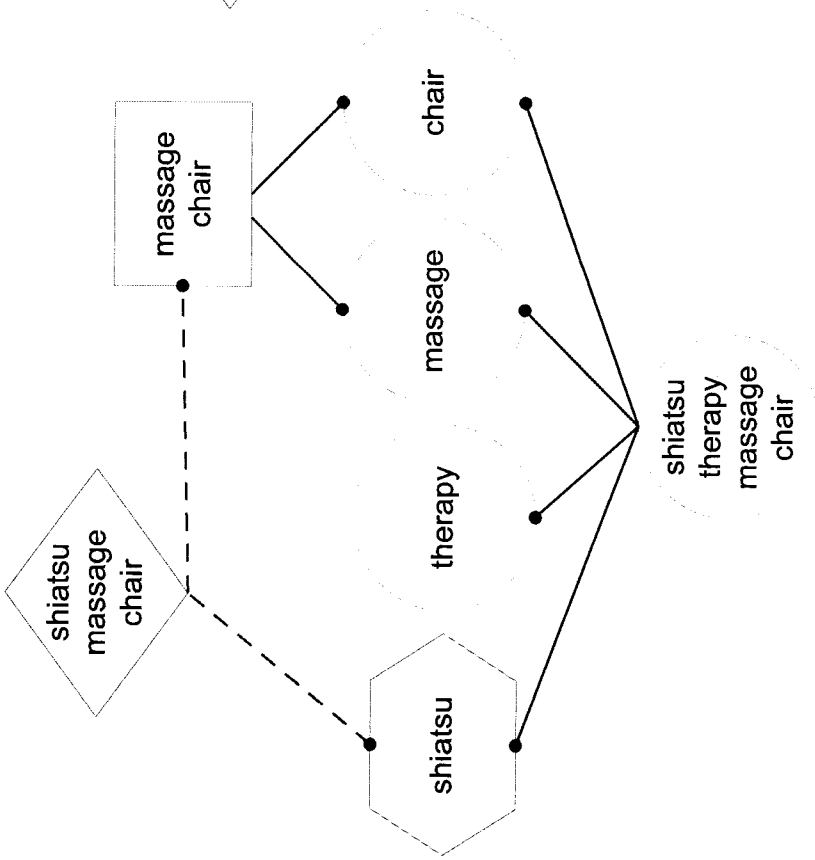
massage chair
chair
massage
therapy
shiatsu therapy massage chair
shiatsu massage chair
shiatsu

FIG. 5E

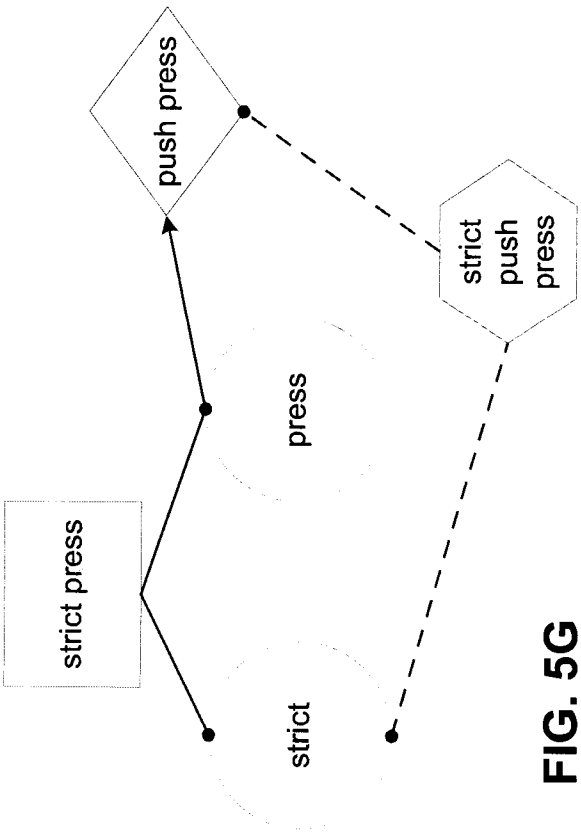
strict press
strict
press
push press
strict push press

FIG. 5G

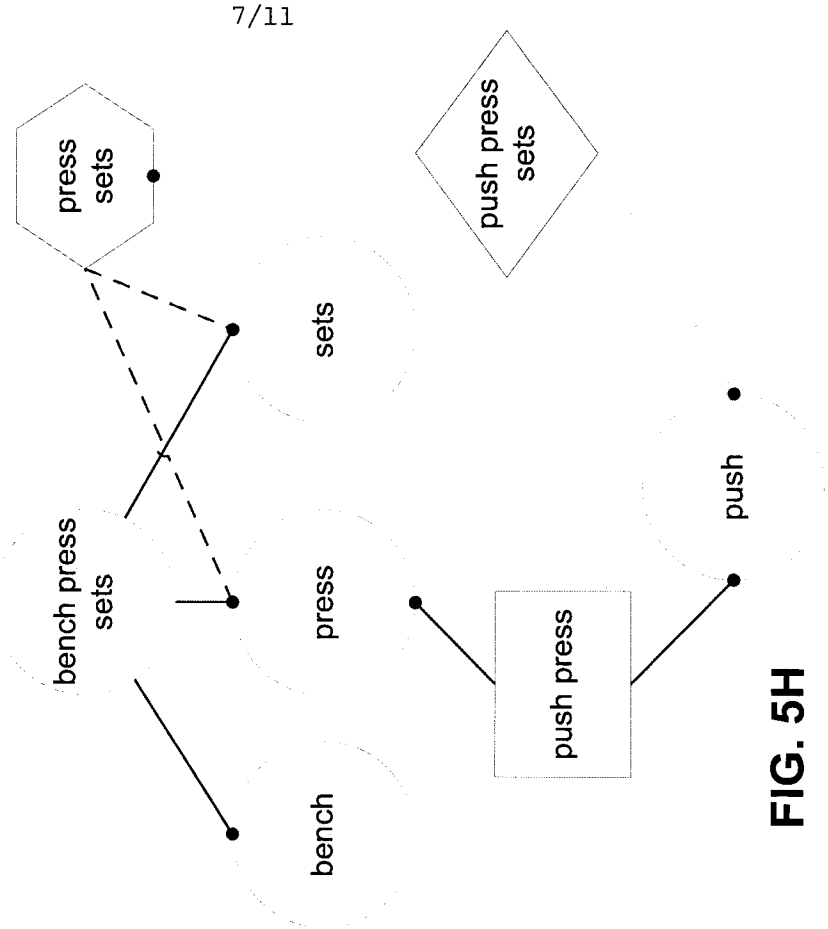
bench press sets
bench
press
sets
press sets
push press
push press sets
push

FIG. 5H

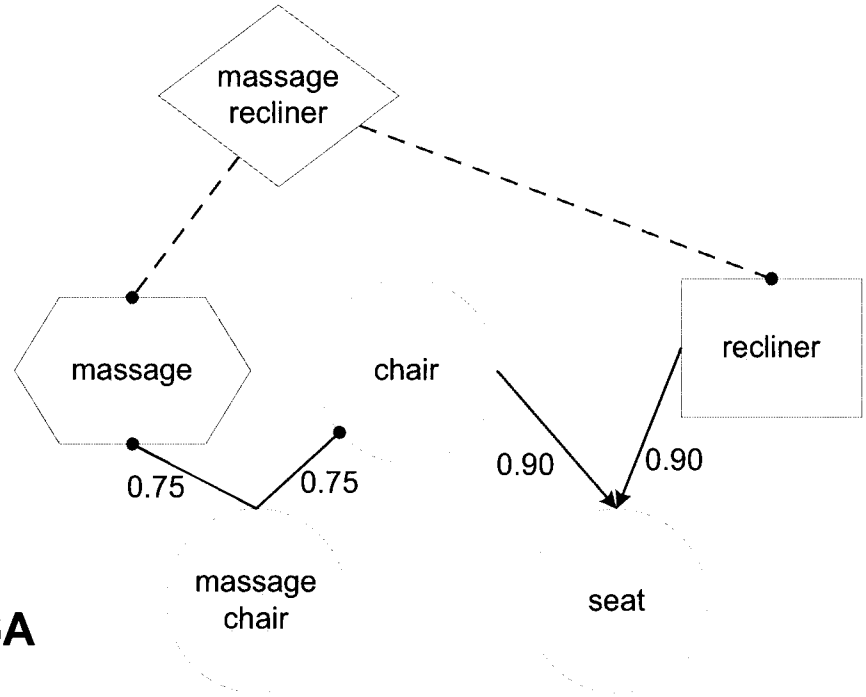
massage recliner
massage
chair
recliner
0.75
0.75
0.90
0.90
massage chair
seat

FIG. 6A

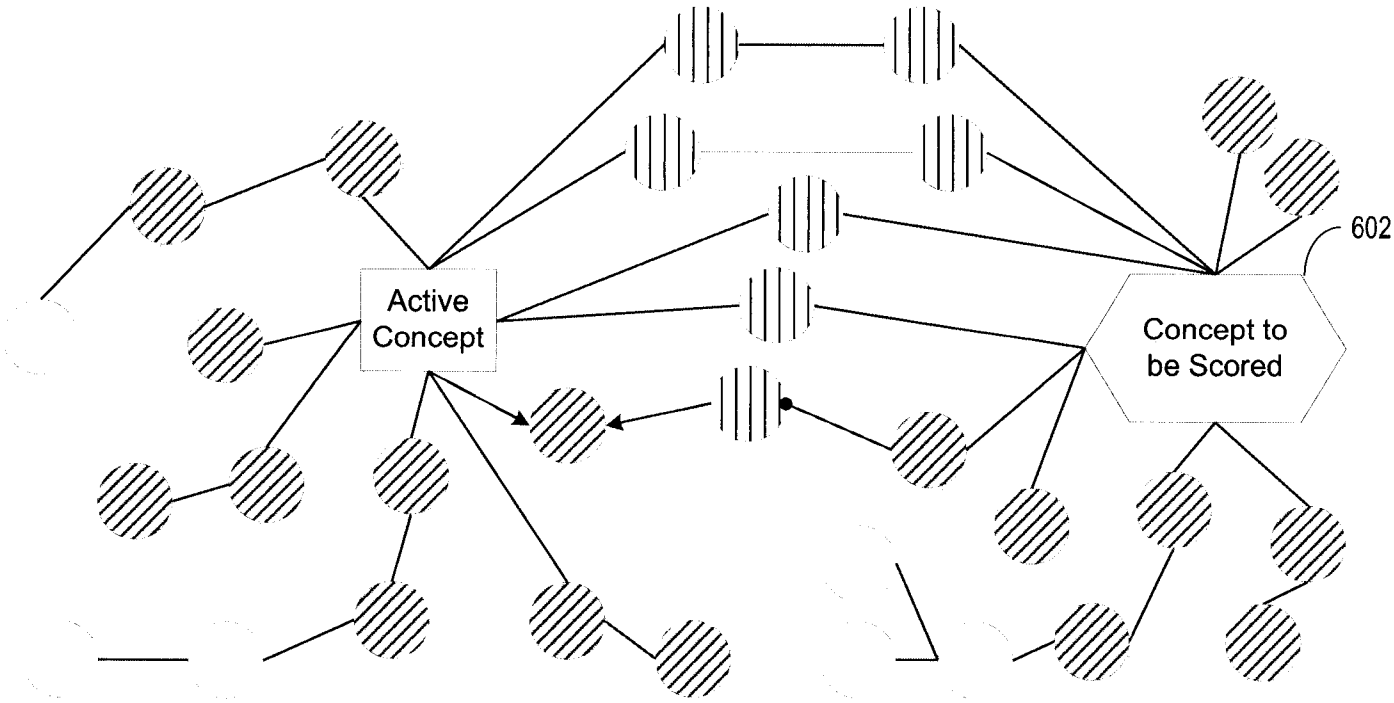
Active Concept
Concept to be Scored
602

FIG. 6B

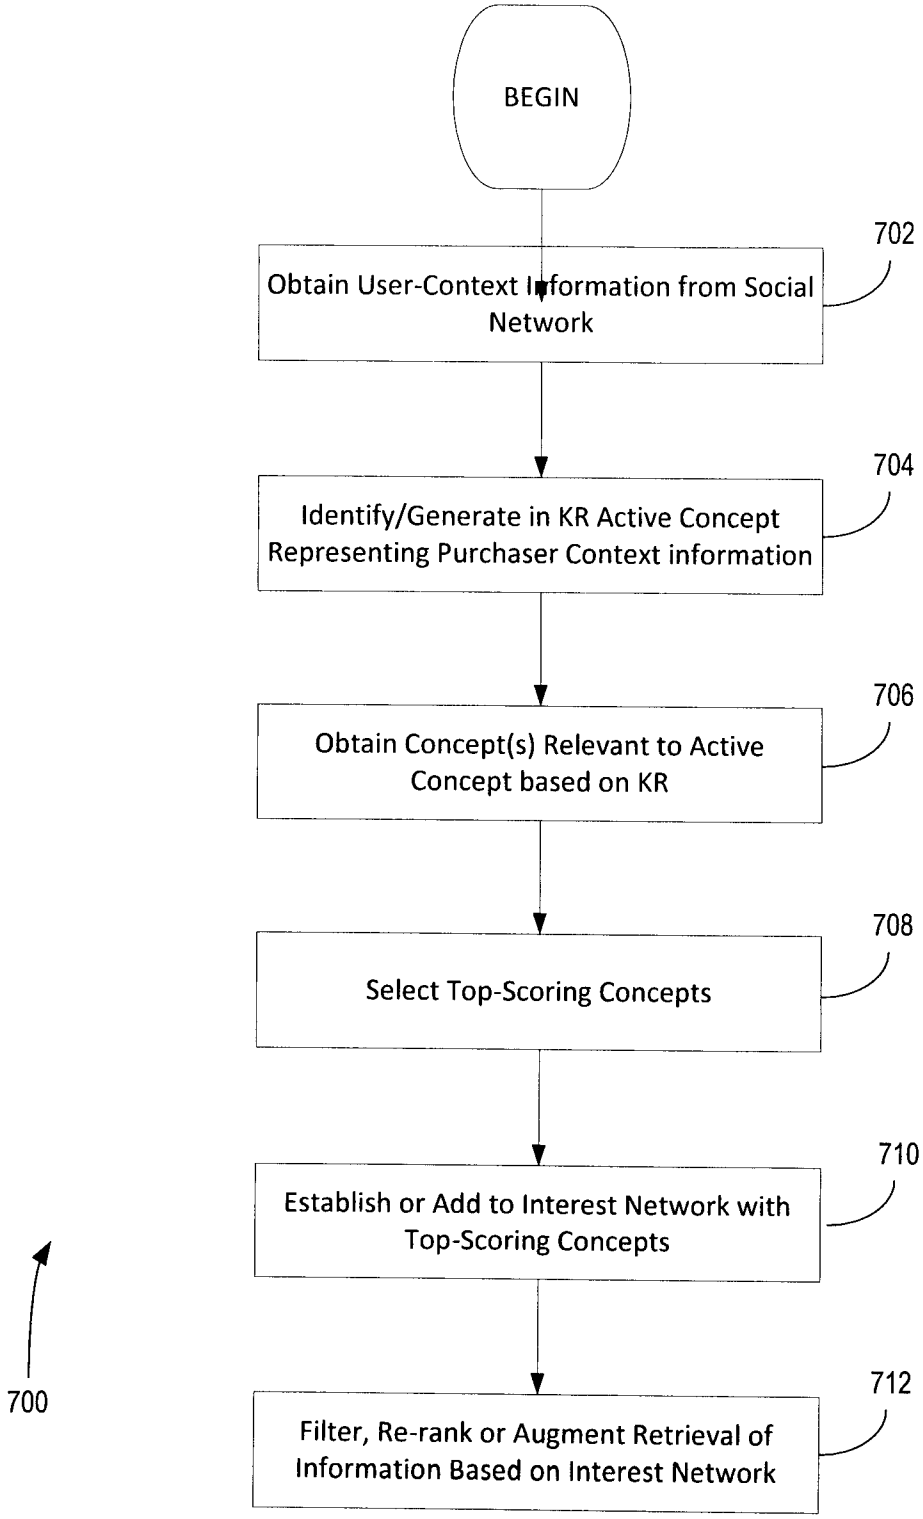
BEGIN
Obtain User-Context Information from Social Network
702
Identify/Generate in KR Active Concept Representing Purchaser Context information
704
Obtain Concept(s) Relevant to Active Concept based on KR
706
Select Top-Scoring Concepts
708
Establish or Add to Interest Network with Top-Scoring Concepts
710
Filter, Re-rank or Augment Retrieval of Information Based on Interest Network
712
700

FIG. 7

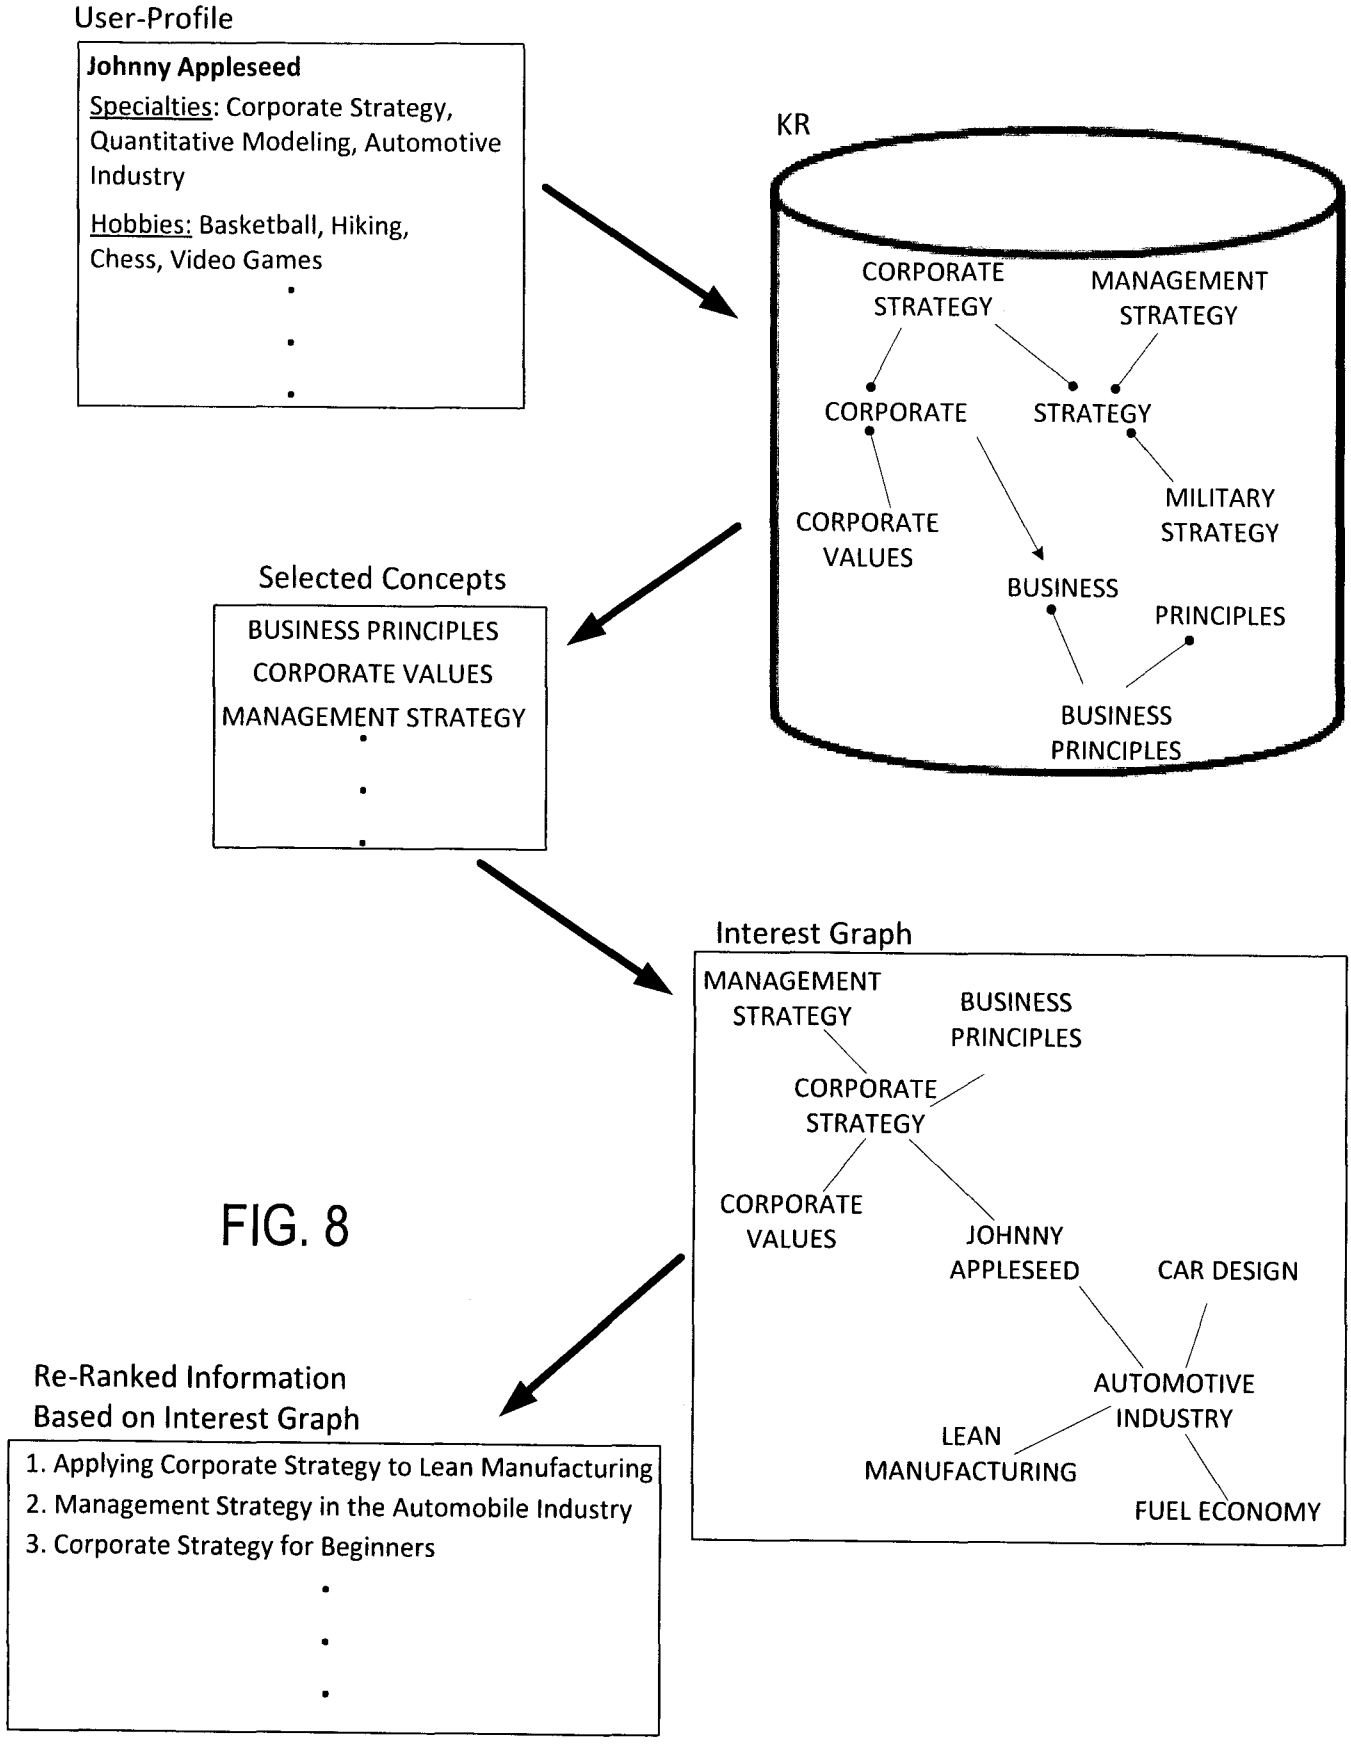
User-Profile
Johnny Appleseed
Specialties: Corporate Strategy, Quantitative Modeling, Automotive Industry
Hobbies: Basketball, Hiking, Chess, Video Games
KR
CORPORATE STRATEGY
MANAGEMENT STRATEGY
CORPORATE
STRATEGY
CORPORATE VALUES
MILITARY STRATEGY
BUSINESS
PRINCIPLES
BUSINESS PRINCIPLES
Selected Concepts
BUSINESS PRINCIPLES
CORPORATE VALUES
MANAGEMENT STRATEGY
Interest Graph
MANAGEMENT STRATEGY
BUSINESS PRINCIPLES
CORPORATE STRATEGY
CORPORATE VALUES
JOHNNY APPLESEED
CAR DESIGN
AUTOMOTIVE INDUSTRY
LEAN MANUFACTURING
FUEL ECONOMY
Re-Ranked Information Based on Interest Graph
1. Applying Corporate Strategy to Lean Manufacturing
2. Management Strategy in the Automobile Industry
3. Corporate Strategy for Beginners

FIG. 8

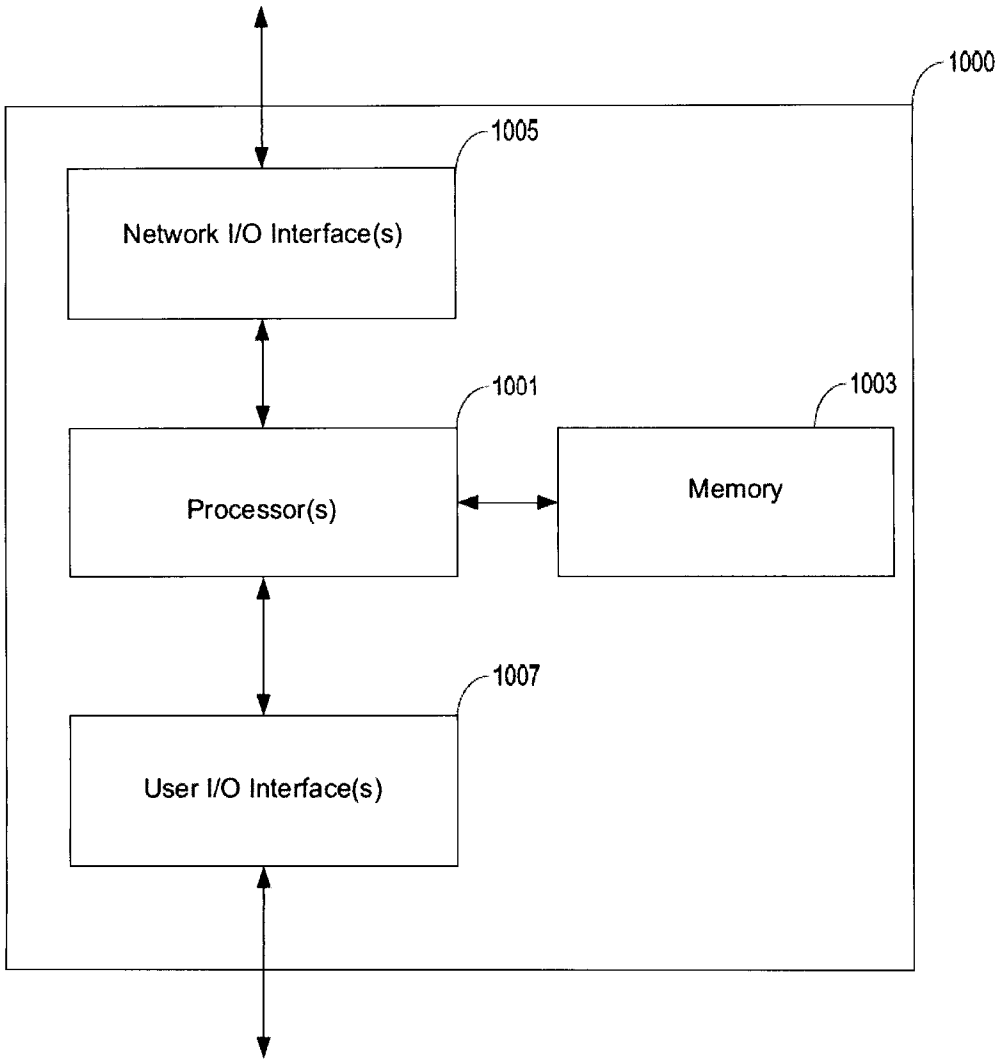
1000
1005
Network I/O Interface(s)
1001
1003
Processor(s)
Memory
1007
User I/O Interface(s)

FIG. 9

INTERNATIONAL SEARCH REPORT

International application No.
PCT/CA2011/001403

A. CLASSIFICATION OF SUBJECT MATTER
IPC: ***H04L 12/16*** (2006.01) , ***G06F 17/30*** (2006.01)
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)
IPCs: H04L-12/16 and G06F-17/30; and
ECLAs: G06Q-10/10, G06Q-30/02, G06F-17/27S, G06F-17/30T2F4, G06F-17/30T4M, G06F-17/30W1, G06F-17/30W1F and H04L-

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

Electronic database(s) consulted during the international search (name of database(s) and, where practicable, search terms used)
Canadian Patent Database, and EPOQUE (EPODOC, English Full Text Databases) - Search terms used: social network, semantic, semantic network, context, concept, ontolog[y,ies], interest network

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2010/149427 A1 (HACID et al.) 29 December 2010 (29-12-2010) Abstract Figure 1	1, 3-5, 7, 9-11, 13 and 15-17

Y	Page 2, lines 1-11, Page 5, line 1 to page 6, line 6 Page 6, lines 7-20, Page 7, lines 19-27, Page 7, line 28 - page 8, line 2 Page 8, lines 3-26, Page 9, line 1 - page 10, line 15 Page 13, line 6 - page 14, line 14 Claims 1-9	2, 6, 8, 12, 14 and 18
Y	US 2010/0250526 A1 (PROCHÁZKA et al.) 30 September 2012 (30-09-2010) Abstract Figures 2-14 Paragraphs [0004-0009], [0046-0053], [0058-0062] Paragraphs [0064-0071] and [0085-0097]	2, 6, 8, 12, 14 and 18

[**X**] Further documents are listed in the continuation of Box C.

[**X**] 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	Date of mailing of the international search report
05 April 2012 (05-04-2012)	23 May 2012 (23.05.2012)

Name and mailing address of the ISA/CA	Authorized officer
Canadian Intellectual Property Office Place du Portage I, C114 - 1st Floor, Box PCT 50 Victoria Street Gatineau, Quebec K1A 0C9 Facsimile No.: 001-819-953-2476	Donald Lefebvre (819) 997-2822

C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
P,Y	US 2011/0173176 A1 (CHRISTENSEN et al.) 14 July 2011 (14-07-2011) Whole document	1-18
A	US 2010/0280860 A1 (ISKOLD et al.) 04 November 2010 (04-11-2010) Whole document	1-18
A	US 2010/0070448 A1 (OMOIGUI) 18 March 2010 (18-03-2010) Whole document	1-18
A	US 2010/0100546 A1 (KOHLER) 22 April 2010 (22-04-2010) Whole document	1-18
A	US 2009/0327417 A1 (CHAKRA) et al. 31 December 2009 (31-12-2009) Whole document	1-18
A	FIKADU et al., "A Framework for Personalized Information Retrieval Model", Conference Proceedings, Second International Conference on Computer and Network Technology (ICCNT), IEEE, Piscataway, NJ, USA, 23 April 2010, Pages 500-505, XP031682297 Whole document	1-18

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
WO2010149427A1	29-12-2010	FR2947358A1	31-12-2010
US2010250526A1	30-09-2010	US2010250526A1	30-09-2010
US2011173176A1	14-07-2011	US2011173176A1	14-07-2011
US2010280860A1	04-11-2010	US2010280860A1	04-11-2010
US2010070448A1	18-03-2010	AU2004213986A1	02-09-2004
		BR0210589A	26-04-2005
		BRPI0407451A	07-02-2006
		CA2555280A1	02-09-2004
		CN1647070A	27-07-2005
		CN1774814A	17-05-2006
		CN100449769C	07-01-2009
		CN1853180A	25-10-2006
		EA008675B1	29-06-2007
		EP1410258A1	21-04-2004
		EP1410258A4	11-07-2007
		EP1597768A1	23-11-2005
		EP1599811A2	30-11-2005
		EP1599811A4	06-02-2008
		EP1749259A2	07-02-2007
		EP1889233A2	20-02-2008
		EP1952266A2	06-08-2008
		EP1952266A4	20-01-2010
		EP1952286A2	06-08-2008
		JP2005514673A	19-05-2005
		JP2006522388A	28-09-2006
		JP2007524985A	30-08-2007
		KR20050098958A	12-10-2005
		KR100812792B1	12-03-2008
		KR20060004909A	16-01-2006
		MXPA03011976A	01-07-2005
		MXPA05008670A	17-11-2005
		SG160209A1	29-04-2010
		TWI289927B	11-11-2007
		US2004161868A1	19-08-2004
		US6946715B2	20-09-2005
		US2004232314A1	25-11-2004
		US7407830B2	05-08-2008
		US2003126136A1	03-07-2003
		US2004230572A1	18-11-2004
		US2006026113A1	02-02-2006
		US2007016563A1	18-01-2007
		US2007038610A1	15-02-2007
		US2007081197A1	12-04-2007
		US2007260580A1	08-11-2007
		US2008016036A1	17-01-2008
		US2008147716A1	19-06-2008
		US2008147788A1	19-06-2008
		US2008162498A1	03-07-2008
		US2008288456A1	20-11-2008
		US2009050944A1	26-02-2009
		US2009254510A1	08-10-2009
		US2011047148A1	24-02-2011
		WO03001413A1	03-01-2003
		WO2004075299A1	02-09-2004
		WO2004075466A2	02-09-2004

Patent Document Cited in Search Report	Publication Date	Patent Family Member(s)	Publication Date
US2010070448A1 (continued)		WO2004075466A3	28-10-2004
		WO2005103883A1	03-11-2005
		WO2005111789A2	24-11-2005
		WO2005111789A3	23-08-2007
		WO2006124952A2	23-11-2006
		WO2006124952A3	16-04-2009
		WO2007044865A2	19-04-2007
		WO2007044865A3	07-05-2009
		WO2007059536A2	24-05-2007
		WO2007059536A3	02-05-2008
		WO2008014499A2	31-01-2008
		WO2008014499A3	03-07-2008
US2010100546A1	22-04-2010	US2010100546A1	22-04-2010
US2009327417A1	31-12-2009	US2009327417A1	31-12-2009
		US7865592B2	04-01-2011