

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

CORRECTED VERSION

(19) World Intellectual Property Organization  
International Bureau



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

WIPO | PCT

(10) International Publication Number  
**WO 2012/088590 A9**

(51) International Patent Classification:	<i>G06N 5/02</i> (2006.01)	<i>G06Q 30/02</i> (2012.01)	61/430,141	5 January 2011 (05.01.2011)	US
(21) International Application Number:		PCT/CA2011/001402	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
(22) International Filing Date:		30 December 2011 (30.12.2011)	(71) Applicant (for all designated States except US):	PRIMAL FUSION INC. [CA/CA]; 7-258 King Street North, Waterloo, Ontario N2J 2Y9 (CA).	
(25) Filing Language:	English		(72) Inventors; and		
(26) Publication Language:	English		(75) Inventors/Applicants (for US only):	SWEENEY, Peter [CA/CA]; 56 Kilbirnie Court, Kitchener, Ontario N2R1B8 (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 N2K1J5 (CA). HUNT, Anne, Jude [US/US]; 825 Cowper Street, Apartment E, Palo Alto California 94301 (US).	
(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			
61/428,598	30 December 2010 (30.12.2010)	US			
61/428,687	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			

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR USING A KNOWLEDGE REPRESENTATION TO PROVIDE INFORMATION BASED ON ENVIRONMENTAL INPUTS




FIG. 1

(57) **Abstract:** There is disclosed a system and method for using a knowledge representation to provide relevant information based on environmental inputs. In an embodiment, the system and method considers environmental information from members in a crowd to generate a pool of interests based on the semantic relevance concepts associated with those interests. The most prominent concepts of interest may then be the basis for presenting content to the crowd as a whole. In another embodiment, environmental inputs and other surrounding inputs are considered as a user context. The concepts may be identified as relevant from the environmental context and used to present information relevant to the user given his or her surroundings.

WO 2012/088590 A9



(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, 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, PL, PG, PI, 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))
- with information concerning incorporation by reference of missing parts and/or elements (Rule 20.6)

(48) **Date of publication of this corrected version:**

20 September 2012

(15) **Information about Correction:**

see Notice of 20 September 2012

## SYSTEM AND METHOD FOR USING A KNOWLEDGE REPRESENTATION TO PROVIDE INFORMATION BASED ON ENVIRONMENTAL INPUTS

### CROSS-REFERENCE TO RELATED APPLICATIONS

**[0001]** The present application 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. More particularly, the teachings disclosed herein relate to the deployment of systems and methods for delivery of promotional content relevant to individuals in a crowd using synthesis of environmental inputs for context matching.

### 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 delivery of promotional content relevant to individuals in a crowd using synthesis of environmental inputs for context matching. In an embodiment, the system and method synthesizes environmental information and other surrounding information about a crowd in real time, or near real time, to determine one or more contexts associated with one or more groups of individuals in the crowd. These environmental inputs and other surrounding inputs are analyzed and synthesized in order to contextualize the surrounding environment and identify one or more contexts (e.g. common themes or interests) amongst a disparate group of individuals in a crowd, such that suggested activities, advertisements and promotional content directed to the crowd will be more dynamic and focussed on the interests of the individuals of the particular crowd in order to capture the attention of a larger number of individuals in the crowd.

**[0009]** Thus, in an aspect, there is provided a computer-implemented method for using a knowledge representation to provide information based on an environmental input, the method comprising: receiving at least one environmental input as a user-context information associated with a user; obtaining at least one concept in the knowledge representation, wherein the at least one concept is obtained based on a semantic relevance of the at least one concept to the user-context information; and based on the at least one concept, providing information to the user; wherein a concept is represented by a data structure storing data associated with a knowledge representation.

**[0010]** In an embodiment, the knowledge representation comprises a semantic network and the at least one concept is represented by a data structure storing data associated with a node in the semantic network.

**[0011]** In another embodiment obtaining the at least one concept comprises obtained the at least one other concept in the knowledge representation based at least in part on the structure of the knowledge representation.

**[0012]** In another aspect, there is provided a system for using a knowledge representation to provide information based on an environmental input, the system adapted to: receive at least

one environmental input as a user-context information associated with a user; obtain at least one concept in the knowledge representation, wherein the at least one concept is obtained based on a semantic relevance of the at least one concept to the user-context information; and based on the at least one concept, provide information to the user; wherein a concept is represented by a data structure storing data associated with a knowledge representation.

**[0013]** 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 provide information based on an environmental input, the computer-readable medium comprising: code for receiving at least one environmental input as a user-context information associated with a user; code for obtaining at least one concept in the knowledge representation, wherein the at least one concept is obtained based on a semantic relevance of the at least one concept to the user-context information; and code for providing information to the user based on the at least one concept; wherein a concept is represented by a data structure storing data associated with a knowledge representation.

**[0014]** 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 phrasology and terminology employed herein are for the purpose of description and should not be regarded as limiting.

## **BRIEF DESCRIPTION OF DRAWINGS**

**[0015]** 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:

**[0016]** 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.

**[0017]** FIG. 2 is a block diagram of an illustrative client/server architecture that may be used to implement some embodiments of the present disclosure.

**[0018]** 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.

**[0019]** FIG. 4A-4C is an illustration of generating an active concept representing user context information, in accordance with some embodiments of the present disclosure.

**[0020]** 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.

**[0021]** FIGS. 6A-6B illustrate techniques for scoring concepts relevant to an active concept, in accordance with some embodiments of the present disclosure.

**[0022]** FIG. 7 illustrates a system and operating environment in accordance with an embodiment.

**[0023]** FIG. 8 illustrates a network environment in accordance with an embodiment.

**[0024]** FIG. 9 illustrates a computing device on which some embodiments of the present disclosure may be implemented.

**[0025]** The foregoing is a non-limiting summary of the invention, which is defined by the attached claims.

## **DETAILED DESCRIPTION**

**[0026]** As noted above, the present disclosure relates to a system and method for delivery of promotional content relevant to individuals in a crowd using synthesis of environmental inputs for context matching. In an embodiment, the system and method synthesizes environmental information and other surrounding information about a crowd in real time, or near real time, to determine one or more contexts associated with one or more groups of individuals in

the crowd. These environmental inputs and other surrounding inputs are analyzed and synthesized in order to contextualize the surrounding environment and identify one or more contexts (e.g. common themes or interests) amongst a disparate group of individuals in a crowd, such that suggested activities, advertisements and promotional content directed to the crowd will be more dynamic and focussed on the interests of the individuals of the particular crowd in order to capture the attention of a larger number of individuals in the crowd.

**[0027]** 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 (such as one or more groups of individuals in a crowd) 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 in the context of a crowd.

**[0028]** 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.

**[0029]** 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.).

**[0030]** 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.

**[0031]** 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.

**[0032]** 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.

**[0033]** 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, etc.) 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. For example, reference numeral 808 identifies a “defined-by” relationship and reference numeral 806 identifies an “is-a” relationship.

**[0034]** 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.

**[0035]** 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.”

**[0036]** 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.

**[0037]** 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.

**[0038]** 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.

**[0039]** 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.

**[0040]** 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.

**[0041]** 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).

**[0042]** 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.

**[0043]** 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.

**[0044]** 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.

**[0045]** 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.

**[0046]** 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.

**[0047]** 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.

**[0048]** 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.

**[0049]** 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.

**[0050]** 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.

**[0051]** 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.

**[0052]** 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.

**[0053]** 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

**[0054]** 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).

**[0055]** 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.

**[0056]** 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.

**[0057]** 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.

**[0058]** 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.

**[0059]** 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.

**[0060]** 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.

**[0061]** 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.

**[0062]** 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**

**[0063]** 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.

**[0064]** 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.

**[0065]** 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.

**[0066]** 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.

**[0067]** 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.

**[0068]** 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.”

[0069] 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.

[0070] 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.

[0071] 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.

[0072] 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.

**[0073]** 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.

**[0074]** 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.

**[0075]** 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.

**[0076]** 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.

**[0077]** 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.

**[0078]** 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.

**[0079]** 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 labelled “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 labelled “Board of Directors” may also be connected via a “defined-by” edge to a node associated with the concept labelled “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.”

**[0080]** 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 labelled “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.

**[0081]** 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.”

**[0082]** 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.

**[0083]** 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.

**[0084]** 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.

**[0085]** 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**

**[0086]** 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.

**[0087]** 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.

**[0088]** 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.

**[0089]** 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).

**[0090]** 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.

**[0091]** 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.

**[0092]** 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.

**[0093]** 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.

**[0094]** 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.”

**[0095]** 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.”

**[0096]** 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.

**[0097]** 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.”

**[0098]** 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.

**[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 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.”

**[00100]** 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.

**[00101]** 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.

**[00102]** 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.”

**[00103]** 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.”

**[00104]** 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.”

**[00105]** 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.

**[00106]** 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.”

**[00107]** 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.

**[00108]** 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.

**[00109]** 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.”

**[00110]** 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**

**[00111]** 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.

**[00112]** 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***

**[00113]** 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.

**[00114]** 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.

**[00115]** 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.

**[00116]** 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.

**[00117]** 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.

**[00118]** 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).

**[00119]** 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.

**[00120]** 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:

$$\begin{aligned}
 S_{gc} &= \text{Edge}_{\text{recliner-seat}} \times \text{Edge}_{\text{seat-chair}} \times \text{Edge}_{\text{chair-massage chair}} \times \text{Edge}_{\text{massage chair-massage}} \\
 &\quad \times \text{Node}_{\text{massage recliner}} \\
 &= \text{Edge}_{\text{is-a}} \times \text{Edge}_{\text{is-a}} \times \text{Edge}_{\text{defined-by}} \times \text{Edge}_{\text{defined-by}} \times \text{Node}_{\text{synthesized}} \\
 &= 0.9 \times 0.9 \times 0.75 \times 0.75 \times 0.25 = 0.1139
 \end{aligned}$$

**[00121]** 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***

**[00122]** 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.

**[00123]** 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.”

**[00124]** 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.”

**[00125]** 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***

**[00126]** 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.

**[00127]** 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:

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

**[00128]** 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.

**[00129]** 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.

**[00130]** 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.

**[00131]** 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.

**[00132]** 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***

**[00133]** 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.

**[00134]** 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.

**[00135]** 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.

**[00136]** 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.

**[00137]** 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.

**[00138]** 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.

**[00139]** 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.

**[00140]** 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.

**[00141]** 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.

**[00142]** 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.

**[00143]** 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.

**[00144]** 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*

**[00145]** 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.”

**[00146]** 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***

**[00147]** 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.

**[00148]** 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.

**[00149]** 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)**

**[00150]** 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.

**[00151]** 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.

**[00152]** 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).”

**[00153]** 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.

**[00154]** 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.

**[00155]** 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.

**[00156]** 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.

**[00157]** 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.

**[00158]** 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.

**[00159]** 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.

**[00160]** 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 relevance of the individual terms.

## **VI. Use Environmental Inputs as Context to Provide Content to Individuals/Crowds**

**[00161]** In recent years, technologies capable of sensing a variety of environmental inputs have become more affordable and widely available. Inexpensive digital point-and-shoot consumer cameras are available with technology designed to locate and focus on a human face, and shoot a photo upon detecting a smile. Motion and body position sensor devices that can detect the motions and physical gestures of the human body have been introduced as relatively inexpensive peripherals for console gaming systems. Various other types of sensors are becoming more widely available on electronic devices such as wireless mobile devices, touch pads and netbooks. As will be explained in more detail below, the present system and method is able to consider environmental inputs and other surrounding information as a context, in order to provide for example, advertisements and promotions, to a user associated with that context.

**[00162]** Advertising via electronic media has seen phenomenal growth over the last few years, growing hand-in-hand with the expansion of the Internet. It has evolved from randomly displayed, passive advertisements to advertisements targeted to specific individuals based on their demographic and psychographic profile. Examples of such technologies are given by US Patent No. 7062466, US Patent No. 7007074 and US Patent Applications 20050216335, 20070038500, 20080243480, 20080228568 and 20090070219. What these approaches have in common is matching a specific individual or consumer group with the media or content that provides the setting for the advertisement.

**[00163]** Advertising via electronic media may be particularly challenging when an advertiser needs to select an appropriate advertisement or promotion to match the disparate interests of individual members of a crowd. As an illustrative example, such a disparate crowd may be found in a public location such as an airport lounge, a food court in a shopping mall, or a public square. In such environments, a common problem is how to capture the attention of a significant number of the individual members of the crowd in order to make promotional content more effective. Given the typically diverse interests in individuals within a crowd, presenting one dimensional, generic promotional content may result in alienating or not receiving the attention of large segments of the audience.

**[00164]** In accordance with one aspect, the present system and method may use one or more methods to, using a knowledge representation, provide information based on environmental and surrounding input by collecting such inputs volunteered by or authorized by individuals or members of a crowd. The knowledge representation may be used, according to some aspects, to perform a semantic operation (as described in more detail above) in order to provide information to the user. Such inputs may be leveraged as a user-context that provides the basis for the semantic operation that may be performed on a computing device operable to perform one or more methods as previously described with respect to FIG. 1 and the description associated with FIG. 1.

**[00165]** In an embodiment, the semantic operations may be integrated with local electronic media such as digital displays, and at least one wireless service within a localized area, such as a Wi-Fi connection made available to people within a certain operating range. As an

illustrative example, in an airport lounge, members of a crowd sitting in front of one or more electronic displays may participate as members of one or more target groups in the crowd for crowd based advertising.

**[00166]** To encourage participation, the members may be offered incentives, such as a free Wi-Fi connection or reduced data transmission rates by their carrier, in exchange for non-identifiable collection of information. In an embodiment, the collected information may be such that it cannot identify any personal information of any individual member of the crowd, such that privacy is maintained.

**[00167]** The present system and method may then collect, in real-time, information about the various websites that individual members of the target crowd are visiting, searches that members of the crowd are performing, or status updates posted by members of the crowd. This information may be considered as part of the context information associated with a user, as described above with respect to FIG. 1. Synthesis or retrieval operation, as described in more detail above, may then be performed on an active concept derived from these user contexts to semantically expand the range of concepts known to be relevant to the crowd. The intersection or most prevalent concepts in the group may be deemed the areas that are likely to appeal to the largest constituency of the crowd.

**[00168]** Content may then be delivered to the crowd based on concepts representing interests with global appeal to the group of people. It should be noted that while in some embodiments, the interests may be the concepts obtained by semantic operations, in other instances the concepts may provide the starting point for further operations that are used to derive such interests, which in turn may be used to identify relevant content. Such content may serve a wide range of purposes as described above, whether merely for entertaining the crowd or for a specific commercial purpose such as generating messaging such as advertising, or influencing the creation of new content by presenting ideas from a domain that embody “looser” associations to encourage lateral thinking by the consumers.

#### ***VI.A Identifying Interests of Target Groups***

**[00169]** Rather than conveying a generic message to a crowd, the present system and method receives environmental and surrounding inputs from one or more individuals in a crowd and generates one or more contexts relevant to one or more target groups or individuals in the crowd. These one or more contexts may form the basis for one or more active concepts representing user context information, each of which may then be used to generate one or more sets of relevant concepts utilizing a knowledge representation as described in more detail above. Generating a number of semantically relevant concepts to a number of individuals in the crowd may present a rich landscape of information for ascertaining interests relevant to the crowd at large. In some embodiments, these concepts may be used to select one or more promotional messages, or other germane content, that may be more relevant to one or more groups of individuals within the crowd.

**[00170]** The ability to identify relevant concepts and interests that may appeal to a large number of users with a broad variety of interests may be a cost-effective approach in presenting advertisements and promotions, as well as suggesting ecommerce-based activities. In an illustrative embodiment, this may involve: i) developing crowd based interests by identifying semantically relevant concepts and/or interest of individuals within group or crowd; ii) lining up advertisers behind each theme or interest; and (iii) flashing advertisement/promotional content to the crowd from a queue of advertisers.

**[00171]** As an illustrative example, referring to FIG. 7, shown is another illustrative embodiment in which environmental context information may be gathered from individuals. Such environmental context information may be gathered in the setting of a crowd according to one aspect. It should be appreciated that the instant invention is not limited to applying a semantic operation to a context ascertained from the environment where there is a crowd. Indeed, the inventors have recognized that the leveraging the multi-modal input as the basis for performing semantic operations and enhancing a user's information retrieval experience may be performed whether the user is with other people or by themselves.

**[00172]** As another illustrative example, in the aspect of the present system and method that may be applied to a crowd, the techniques and methodologies may be employed to identify an interest applicable to crowds of mass scale, such as those that may be found at a sports

stadium. For example, a crowd in a hockey game may be exposed to various digital advertisement banners over the course of an evening attending a hockey game. In accordance with the present system and method, information from individuals within large audience may be semantically expanded to produce a large pool of relevant concepts, the intersection or most prevalent of which may be identified. Prevalence, as applied to the crowd at large, may be determined using any of the scoring techniques described above. This information may then be used to dynamically customize the advertisements being presented to the crowd in real time. Thus, any shift in the trending interests may be immediately addressed by corresponding changes to the advertisements.

**[00173]** The present system and method may be adapted to collect information about members of a target audience within a threshold distance of a predetermined, fixed location. For example, a threshold distance may be determined by the operating distance to connect to a Wi-Fi connection. In another aspect, the threshold distance may be a function of the visual and/or audio access that an advertising medium may reach, ensuring consideration of contexts from only those audience members who will benefit from the advertising.

**[00174]** Alternatively, an electronic display or advertising medium may be provided on moving platform, such as a billboard truck, that may move around from location to location to target a particular crowd. For example, a crowd assembled at an outdoor concert in a town square, or a street festival event may be targeted by such a platform. In another embodiment, rather than targeting members of a crowd who are largely stationary (e.g. in a crowd in an airport lounge, or a crowd sitting in a sports stadium), the present system and method may also be adapted to gather information on members who are collectively moving, for example drivers and passengers in automobiles approaching an electronic advertisement billboard at the side of a road, an electronic banner scrolling across an overhead bridge, etc. This may also apply to pedestrians that are walking, such as a thoroughfare that may have high foot-traffic. By collecting information about the members of the crowd approaching the electronic display or banner, the system and method of the present invention may be adapted to modify the relevant content or advertisements on the electronic billboards or banners.

**[00175]** The individual user-context within a crowd in any of the foregoing situations may be identified based on the approaches illustrated with respect to acts 102 of FIG. 1 and other teachings for identifying a user-context as described above. Similarly, obtaining semantically relevant concepts to each user may be accomplished by any number of suitable approaches, including but not limited to those outlined above in acts 104-108 and other methodologies as described above. In addition to the techniques and approaches for gathering context as described above, environmental input may provide a user-context as detailed below.

#### ***VI.B Receiving Environment Based Inputs***

**[00176]** As previously noted, the present disclosure also relates to a system and method for identifying interest based on concepts deemed relevant to an individual from a semantic operation performed on the user's environmental context. The present system and method may be practiced in various embodiments. A suitably configured computer device, and associated communications networks, devices, software and firmware may provide a platform for enabling one or more embodiments as described above. An illustrative computer device and an embodiment of the system is shown in FIGS. 7-9 as described below.

**[00177]** As shown in FIGS. 7 and 8, the system and method of the present disclosure may be practiced within a wireless network operating environment which may include mobile devices, or other types of wireless technology which may individually or collectively identify members in a crowd. For example, and not by way of limitation, such mobile devices or wireless technology may include featured mobile phones or smart phones, touch pads, net books, laptops, or any other type of device incorporating wireless technology and adapted to capture and store user-context information. The mobile devices and wireless technology may be based on 3G technologies, but may also be based on 4G and 5G technologies that will enable even faster collection and analysis about users. Other types of technologies that may be used include, for example, radio-frequency identification (RFID) technology that may be carried or worn by various members in the target crowd.

**[00178]** The present system and method may also be implemented as a computer-readable/useable medium that includes computer program code to enable one or more computer devices to implement each of the various process steps in a method in accordance with the

present invention. It is understood that the terms computer-readable medium or computer useable medium comprises one or more of any type of physical embodiment of the program code. In particular, the computer-readable/useable medium can comprise program code embodied on one or more portable storage articles of manufacture (e.g. an optical disc, a magnetic disk, a tape, etc.), on one or more data storage portioned of a computing device, such as memory associated with a computer and/or a storage system.

**[00179]** Many different types of sensing devices may be used to detect and analyze environmental inputs including, and not by way of limitation, image and video input sensors, acoustic input sensors, touch and pressure sensors, motion and orientation sensors, global positioning and speed sensors, temperature and humidity sensors, electric/magnetic field sensors, vapour and chemical sensors, and the like. Each, all or any suitable combination of these inputs may be translated into a label that may be the basis for identifying or generating an active concept as described above with respect to act 104 in process 100 or process 300 along with associated text.

**[00180]** For example, an acoustic input sensor may be applied to detect ambient music. Any of a number of available commercial applications, especially popular on smart phones and tablet devices, may then be used to identify the song, artist genre of music. The genre of music may for example then be used as the context for generating the active concept. As another non-limiting example, categorizations may be assigned to weather conditions, so that when humidity and temperature detected by sensors exceed certain levels, a context of “hot and humid” may be provided as the user-context. It may be envisioned that certain products (e.g. fans) or literature on certain topics (e.g. body hydration) may be of interest to individuals in this setting. Such product advertisements and article suggestions may be ascertained from relevant interests and concepts (e.g. “body hydration”) derived from semantic operations performed on an active concept identified from the user context (e.g. “hot and humid”).

**[00181]** The various types of sensors may be connected to one or more analyzers for analyzing the collected inputs. For example, if the input is from an image sensor, analysis may include determination of light intensity and color to determine ambient mood, or the use of facial recognition technology to determine if a subject is smiling, and even infer whether a subject is

male or female. Images can be processed by utilizing computer animation technologies, or by comparing to a repository of images and identifying a whole or a part of an image from one or more images stored in a database. Image can be captured from high resolution networked cameras or cameras based in the frame of the display.

**[00182]** Similarly, video sensors may be used to detect motion and simple gestures, and compare them to a motion or gesture repository in order to find a match. Video input may also be a video that is being displayed or broadcast within the local environment on a LAN or internet. Any or all of these video inputs may then be processed by associated content/meta tags and/or speech, images or text within the video.

**[00183]** Environmental input may include other types of audio input, such as speech or other acoustic inputs. The audio inputs may be generally classified as object generated or human generated. Object generated inputs may provide a way to capture multimedia inputs without interfering with an individual user's privacy. Some examples of object generated inputs include: laughter, music, public announcements or other publicly available inputs that may not affect the privacy of the user.

**[00184]** In an embodiment, the audio input may distinguish between dialects when the user's voice is considered as an environmental input, and this may be used to build a profile of the user based on their dialect. For example, a dialect identified as being of Cajun origins may be used to identify or generate an active concept "Cajun." A semantically relevant concept, using the techniques as described above with respect to Fig. 1, may then obtain "French Cuisine" as a relevant concept to the user, from which advertisements for local dining establishments to that effect may be suggested. Speech may also be converted into text and the text may be analyzed for semantic meaning or translated into different languages as may be appropriate.

**[00185]** In addition to speech and language, other audio input indicative of emotion or mood may also be captured and processed for further synthesis. For example, the analysis may include determination of volume and intensity of a user's voice to determine the mood of the user, or the level of ambient noise to determine if the user is in the middle of a crowd, or sitting in a quiet room.

**[00186]** Another type of sensor may receive motion input to determine how active the user currently is – whether moving around rapidly during the middle of a busy weekday, or relaxing comfortably at home on a weekend (this context may be considered along with temporal inputs such as the actual time of day and day of the week, which may further suggest in which type of activity the user may be engaged). For example, constant fidgeting may suggest a context of uneasiness that may in turn be used to identify or generate “anxiousness” or “restlessness” as an active concept. The approaches, outlined above for example in Fig. 1, may be used to generate semantically relevant concept that may be of interest to the user, such as “relaxation techniques.” Such concepts of interest may be employed to retrieve and present content, such as articles on relaxation techniques, to the user.

**[00187]** As described above, one or a single of sensors designed to capture environmental inputs may be built in to an electronic device. In an embodiment, the general computing device may be a mobile device, such as a smart phone or touch pad, for example. The mobile device may include a gestural input, accelerometers, a GPS for location data, picture and sound (or music) sensors, a built-in camera, and other types of sensors for receiving input from the surrounding environment. These different types of inputs may be collected in different ways, and may require participation from the user to collect the input.

**[00188]** If multiple inputs are received by the sensors simultaneously, then the multiple inputs from all types of sensors may be considered to more accurately measure the surrounding environment. More generally, multi-media and motion sensor inputs may be received and analyzed in real time, or near real time, and the analyzed data may be aggregated in order to better assess a user’s surrounding environment at any particular moment. By more comprehensively capturing a user’s environmental context at any given time, a greater number and higher quality of relevant concepts and interests may be obtained, for example, to a suggested activity, advertisement, or promotion matching the user’s surrounding.

**[00189]** In another aspect, machine vision sensors and motion sensors may be used to determine a motion or orientation of a user to determine a user’s activity or mood. For example, the sensors may determine what type of motion or gesture is occurring by comparing the motion to a databank of pre-recorded motions. By analyzing the motion by making such a comparison,

the machine vision or motion sensor analyzers can translate the motion into meaningful data. For example, detection of a feature on a user's face may be translated into a text or graphic equivalent (i.e. if a computing device receives data suggesting a user is smiling or in a particular emotional state, a linguistic recitation of that state may provide a context, that in turn may be used to generate semantically relevant concepts that are of particular interest to the user given his or her current emotional state).

**[00190]** When many different motions or gestures are occurring at the same time or in quick succession, the input data may be analyzed to determine a predominant type of motion, as long as it can be identified. Alternatively, the different types of gestures or motions may be considered in sequence, and the sequence may be used to consider what types of changes may be occurring based on signatures or patterns within the data.

**[00191]** In another aspect, the present system and method may also utilize sensor-based capabilities embedded or affixed to various articles that may be worn by a user. For example, sensors may be built into a shoe, or worn in some article of clothing. For example, sensors built into clothing or shoes may transmit a signal that may be detected by wireless detection means. Such sensors may be connected to user's mobile devices via short-range communication protocols such as Bluetooth or IrDA. Thus, the present system and method may receive inputs from multiple different types of sensors, such that a more comprehensive profile of a user and his/her environment may be created.

**[00192]** Any suitable computing device may be used which allows collection of inputs to be processed locally on the device that is directly coupled to the sensors picking up environmental inputs. Alternatively, rather than processing inputs locally, some or all of the inputs may be collected locally by the computer device, but the data may be transmitted to a remote location for processing as shown in FIG. 2.

**[00193]** It may be appreciated that the various foregoing environmental inputs are intended as illustrative and not exhaustive. Any suitable environmental input may be considered to capture a user-context, based on linguistic expressions associated with any of the multi-modal data derived from the environmental input. Such an environmental input may be the basis of user-context consistent the approaches described above with respect to acts 102 of FIG. 1 and

other teachings for identifying a user-context as detailed in this disclosure. Similarly, obtaining semantically relevant concepts to each user may be accomplished by any number of suitable approaches, including but not limited to those outlined above in acts 104-108 and other methodologies as explained herein. Accordingly, the present system and method provide a means for considering the user's environment in order to enrich the type of information and/or content that user experiences.

## **VII. Additional Implementation Detail**

**[00194]** 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 of an illustrative computing device 1000 that may be used to implement any of the above-discussed computing devices.

**[00195]** 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.

**[00196]** 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.

**[00197]** 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.

**[00198]** 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.

**[00199]** 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.

**[00200]** 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.

**[00201]** 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.

**[00202]** 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.

**[00203]** 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.

**[00204]** 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.

**[00205]** Also, various inventive concepts may be embodied as one or more methods, of which multiple examples have been provided (e.g., processes 100, 300). 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.

**[00206]** 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.

**[00207]** The indefinite articles “a” and “an,” as used herein, unless clearly indicated to the contrary, should be understood to mean “at least one.”

**[00208]** 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.

**[00209]** 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.

**[00210]** 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.

**[00211]** 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.

**[00212]** 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 using a knowledge representation to provide information based on an environmental input, the method comprising:
  - receiving at least one environmental input as a user-context information associated with a user;
  - obtaining at least one concept in the knowledge representation, wherein the at least one concept is obtained based on a semantic relevance of the at least one concept to the user-context information; and
  - based on the at least one concept, providing information to the user, wherein a concept is represented by a data structure storing data associated with a knowledge representation.
2. The computer-implemented method of claim 1, wherein the knowledge representation comprises a semantic network and the at least one concept is represented by a data structure storing data associated with a node in the semantic network.
3. The computer-implemented method of claim 1, wherein obtaining the at least one concept comprises obtaining the at least one other concept in the knowledge representation based at least in part on the structure of the knowledge representation.
4. The computer-implemented method of claim 1, further comprising utilizing one or more sensors to sense the at least one environmental inputs for the user.
5. The computer-implemented method of claim 4, wherein the one or more sensors comprise one or more in-range wireless devices associated with the user.
6. The computer-implemented method of claim 5, wherein the wireless devices comprise one or more of mobile phones, smart phones, touch pads, net books, laptops, or any other wireless device incorporating wireless technology and adapted to store information about one or more individuals in the crowd.

7. The computer-implemented method of claim 1, wherein the environmental input comprises at least one of audio data, video data, light data, geo-location data and/or motion data.

8. The computer-implemented method of claim 1, wherein the information comprises advertising or promotional content.

9. A system for using a knowledge representation to provide information based on an environmental input, the system adapted to:

receive at least one environmental input as a user-context information associated with a user;

obtain at least one concept in the knowledge representation, wherein the at least one concept is obtained based on a semantic relevance of the at least one concept to the user-context information; and

based on the at least one concept, provide information to the user;

wherein a concept is represented by a data structure storing data associated with a knowledge representation.

10. The system of claim 9, wherein the knowledge representation comprises a semantic network and the at least one concept is represented by a data structure storing data associated with a node in the semantic network.

11. The system of claim 9, wherein the system is further adapted to obtain the at least one other concept in the knowledge representation based at least in part on the structure of the knowledge representation.

12. The system of claim 9, wherein the system is further adapted to utilize one or more sensors to sense the at least one environmental inputs for the user.

13. The system of claim 12, wherein the one or more sensors comprise one or more in-range wireless devices associated with the user.

14. The system of claim 13, wherein the wireless devices comprise one or more of mobile phones, smart phones, touch pads, net books, laptops, or any other wireless device incorporating wireless technology and adapted to store information about one or more individuals in the crowd.

15. The system of claim 9, wherein the environmental input comprises at least one of audio data, video data, light data, geo-location data and/or motion data.

16. The system of claim 9, wherein the information comprises advertising or promotional content.

17. A non-transitory computer-readable medium storing computer code that when executed on a computer device adapts the device to provide information based on an environmental input, the computer-readable medium comprising:

code for receiving at least one environmental input as a user-context information associated with a user;

code for obtaining at least one concept in the knowledge representation, wherein the at least one concept is obtained based on a semantic relevance of the at least one concept to the user-context information; and

code for providing information to the user based on the at least one concept;

wherein a concept is represented by a data structure storing data associated with a knowledge representation.

18. The non-transitory computer-readable medium of claim 17, wherein the knowledge representation comprises a semantic network and the at least one concept is represented by a data structure storing data associated with a node in the semantic network.

19. The non-transitory computer-readable medium of claim 17, wherein the code for obtaining the at least one concept comprises instructions for obtaining the at least one other concept in the knowledge representation based at least in part on the structure of the knowledge representation.

20. The non-transitory computer-readable medium of claim 17, further comprising code for utilizing one or more sensors to sense the at least one environmental inputs for the user.

21. A computer-implemented method for providing information selected from a large set of digital content to at least one user, the method comprising:

receiving user context information associated with the at least one user;

identifying or generating, using at least one processor executing stored program instructions, a first concept in a semantic network, the first concept representing at least a portion of the user context information;

obtaining at least one concept, including a second concept, semantically relevant to the first concept at least in part, by synthesizing the second concept based on the first concept and at least one other concept in the semantic network; and

providing information to the at least one user, wherein the information is selected by using the first concept and the at least one obtained concept semantically relevant to the first concept,

wherein the first concept in the semantic network is represented by a data structure storing data associated with a node in the semantic network.

22. The computer-implemented method of claim 21, wherein synthesizing the second concept comprises:

identifying the at least one other concept in the semantic network based at least in part on the structure of the semantic network; and

synthesizing the second concept from the first concept and the at least one other concept.

23. The computer-implemented method of claim 21, wherein synthesizing the second concept comprises using an addition operation based on an attribute co-definition technique.

24. The computer-implemented method of claim 21, wherein synthesizing the second concept comprises using an addition operation based on an analogy-by-parent technique and/or an analogy-by-sibling technique.

25. The computer-implemented method of claim 21, wherein synthesizing the second concept comprises using an addition operation based on an attribute commonality technique.

26. The computer-implemented method of claim 21, wherein synthesizing the second concept comprises using a substitution operation, wherein the substitution operation comprises using a retrieval operation and/or an addition operation.

27. The computer-implemented method of claim 21, wherein obtaining the at least one concept comprises:

obtaining a plurality of concepts semantically relevant to the first concept;

computing a score for one or more concepts in the plurality of concepts, wherein the score for a specific concept is indicative of the semantic relevance of the specific concept to the first concept; and

selecting the at least one concept based on the scores computed for the one or more concepts.

28. The computer-implemented method of claim 27, wherein computing a score for a concept comprises using at least one measure of relevance from among generation certainty, concept productivity, Jaccard, statistical coherence, and/or cosine similarity.

29. The computer-implemented method of claim 21, wherein identifying or generating the first concept comprises:

determining whether the portion of the user context information matches an identifier of a concept in the semantic network; and

when it is determined that the at least a portion of the user context information does not match an identifier of a concept in the semantic network, generating the first concept in the semantic network.

30. The computer-implemented method of claim 21, wherein generating the first concept in the semantic network comprises identifying a concept in the semantic network covering more

words in the portion of the user context information than any other concept in the semantic network.

31. The computer-implemented method of claim 21, wherein the user-specific context information comprises at least one of a search query provided by the user, demographic information associated with the user, information from the user's browsing history, information typed by the user, and/or information highlighted by the user.

32. The computer-implemented method of claim 21, wherein the semantic network is represented by a data structure embodying a directed graph comprising a plurality of nodes and a plurality of edges, wherein each node is associated with a concept and an edge between two nodes represents a relationship between the two corresponding concepts.

33. The computer-implemented method of claim 21, wherein the information comprises one or more advertisements and/or one or more product recommendations from one or more other users.

34. The computer-implemented method of claim 21, wherein the information comprises content appearing on or accessible through a website.

35. The computer-implemented method of claim 21, wherein providing information to the at least one user comprises:

creating a search query that includes terms from the first concept and the at least one obtained concept; and

providing the user with information associated with the search results obtained based on the search query.

36. A system for providing information selected from a large set of digital content to at least one user, the system comprising:

at least one processor configured to perform a method comprising receiving user context information associated with the at least one user;

identifying or generating a first concept in a semantic network, the first concept representing at least a portion of the user context information;

obtaining at least one concept, including a second concept, semantically relevant to the first concept at least in part, by synthesizing the second concept based on the first concept and at least one other concept in the semantic network; and

providing information to the at least one user, wherein the information is selected by using the first concept and the at least one obtained concept semantically relevant to the first concept,

wherein the first concept in the semantic network is represented by a data structure storing data associated with a node in the semantic network.

37. The system of claim 36, wherein synthesizing the second concept comprises:

identifying the at least one other concept in the semantic network based at least in part on the structure of the semantic network; and

synthesizing the second concept from the first concept and the at least one other concept.

38. The system of claim 36, wherein synthesizing the second concept comprises using an addition operation based on an attribute co-definition technique.

39. The system of claim 36, wherein synthesizing the second concept comprises using an addition operation based on an analogy-by-parent technique and/or an analogy-by-sibling technique.

40. The system of claim 36, wherein synthesizing the second concept comprises using an addition operation based on an attribute commonality technique.

41. The system of claim 36, wherein synthesizing the second concept comprises using a substitution operation, wherein the substitution operation comprises using a retrieval operation and/or an addition operation.

42. The system of claim 36, wherein obtaining the at least one concept comprises:

obtaining a plurality of concepts semantically relevant to the first concept;

computing a score for one or more concepts in the plurality of concepts, wherein the score for a specific concept is indicative of the semantic relevance of the specific concept to the first concept; and

selecting the at least one concept based on the scores computed for the one or more concepts.

43. The system of claim 42, wherein computing a score for a concept comprises using at least one measure of relevance from among generation certainty, concept productivity, Jaccard, statistical coherence, and/or cosine similarity.

44. The system of claim 36, wherein identifying or generating the first concept comprises:

determining whether the portion of the user context information matches an identifier of a concept in the semantic network; and

when it is determined that the at least a portion of the user context information does not match an identifier of a concept in the semantic network, generating the first concept in the semantic network.

45. The system of claim 36, wherein generating the first concept in the semantic network comprises identifying a concept in the semantic network covering more words in the portion of the user context information than any other concept in the semantic network.

46. The system of claim 36, wherein the user-specific context information comprises at least one of a search query provided by the user, demographic information associated with the user, information from the user's browsing history, information typed by the user, and/or information highlighted by the user.

47. The system of claim 36, wherein the semantic network is represented by a data structure embodying a directed graph comprising a plurality of nodes and a plurality of edges, wherein each node is associated with a concept and an edge between two nodes represents a relationship between the two corresponding concepts.

48. The system of claim 36, wherein the information comprises one or more advertisements and/or one or more product recommendations from one or more other users.

49. The system of claim 36, wherein the information comprises content appearing on or accessible through a website.

50. The system of claim 36, wherein providing information to the at least one user comprises:  
creating a search query that includes terms from the first concept and the at least one obtained concept; and

providing the user with information associated with the search results obtained based on the search query.

51. At least one non-transitory computer readable storage medium storing processor-executable instructions that when executed by at least one processor, cause the at least one processor to perform a method for providing information selected from a large set of digital content to at least one user, the method comprising:

receiving user context information associated with the at least one user;

identifying or generating a first concept in a semantic network, the first concept representing at least a portion of the user context information;

obtaining at least one concept, including a second concept, semantically relevant to the first concept at least in part, by synthesizing the second concept based on the first concept and at least one other concept in the semantic network; and

providing information to the at least one user, wherein the information is selected by using the first concept and the at least one obtained concept semantically relevant to the first concept,

wherein the first concept in the semantic network is represented by a data structure storing data associated with a node in the semantic network.

52. The at least one non-transitory computer readable storage medium of claim 31, wherein synthesizing the second concept comprises:

identifying the at least one other concept in the semantic network based at least in part on the structure of the semantic network; and

synthesizing the second concept from the first concept and the at least one other concept.

53. The at least one non-transitory computer readable storage medium of claim 51, wherein synthesizing the second concept comprises using an addition operation based on an attribute co-definition technique.

54. The at least one non-transitory computer readable storage medium of claim 51, wherein synthesizing the second concept comprises using an addition operation based on an analogy-by-parent technique and/or an analogy-by-sibling technique.

55. The at least one non-transitory computer readable storage medium of claim 51, wherein synthesizing the second concept comprises using an addition operation based on an attribute commonality technique.

56. The at least one non-transitory computer readable storage medium of claim 51, wherein synthesizing the second concept comprises using a substitution operation, wherein the substitution operation comprises using a retrieval operation and/or an addition operation.

57. The at least one non-transitory computer readable storage medium of claim 51, wherein obtaining

the at least one concept comprises:

obtaining a plurality of concepts semantically relevant to the first concept;

computing a score for one or more concepts in the plurality of concepts, wherein the score for a specific concept is indicative of the semantic relevance of the specific concept to the first concept; and

selecting the at least one concept based on the scores computed for the one or more concepts.

58. The at least one non-transitory computer readable storage medium of claim 57, wherein computing a score for a concept comprises using at least one measure of relevance from among generation certainty, concept productivity, Jaccard, statistical coherence, and/or cosine similarity.

59. The at least one non-transitory computer readable storage medium of claim 51, wherein identifying or generating the first concept comprises:

determining whether the portion of the user context information matches an identifier of a concept in the semantic network; and

when it is determined that the at least a portion of the user context information does not match an identifier of a concept in the semantic network, generating the first concept in the semantic network.

60. The at least one non-transitory computer readable storage medium of claim 51, wherein generating the first concept in the semantic network comprises identifying a concept in the semantic network covering more words in the portion of the user context information than any other concept in the semantic network.

61. The at least one non-transitory computer readable storage medium of claim 51, wherein the user-specific context information comprises at least one of a search query provided by the user, demographic information associated with the user, information from the user's browsing history, information typed by the user, and/or information highlighted by the user.

62. The at least one non-transitory computer readable storage medium of claim 51, wherein the semantic network is represented by a data structure embodying a directed graph comprising a plurality of nodes and a plurality of edges, wherein each node is associated with a concept and an edge between two nodes represents a relationship between the two corresponding concepts.


63. The at least one non-transitory computer readable storage medium of claim 51, wherein the information comprises one or more advertisements and/or one or more product recommendations from one or more other users.


64. The at least one non-transitory computer readable storage medium of claim 51, wherein the information comprises content appearing on or accessible through a website.


65. The at least one non-transitory computer readable storage medium of claim 51, wherein providing information to the at least one user comprises:

creating a search query that includes terms from the first concept and the at least one obtained concept; and




providing the user with information associated with the search results obtained based on the search query.

**FIG. 1**

**FIG. 2**

**FIG. 3**

4/11

**FIG. 4A****FIG. 4B****FIG. 4C**




FIG. 5A




FIG. 5B




FIG. 5C






FIG. 5D

**FIG. 5F****FIG. 5E**

7/11




FIG. 5H




FIG. 5G





FIG. 6A



**FIG. 6B**

9/11




FIG. 7





FIG. 8

11/11

**FIG. 9**