



(51) International Patent Classification:

G06Q 50/00 (2006.01) G06F 16/00 (2019.01)
G06N 99/00 (2010.01)

(21) International Application Number:

PCT/US2018/039268

(22) International Filing Date:

25 June 2018 (25.06.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/660,876 20 April 2018 (20.04.2018) US
15/967,239 30 April 2018 (30.04.2018) US

(71) Applicant: **FACEBOOK, INC.** [US/US]; 1601 Willow Road, Menlo Park, CA 94025 (US).

(72) Inventors: **NATARAJAN, Vivek**; 1601 Willow Road, Menlo Park, CA 94025 (US). **YANG, Wenhai**; 1601 Willow Road, Menlo Park, CA 94025 (US). **LIU, Honglei**; 1601 Willow Road, Menlo Park, CA 94025 (US). **KUMAR, Anuj**; 1601 Willow Road, Menlo Park, CA 94025 (US).

(74) Agent: **EVERY, Matthew**; Baker Botts L.L.P., 2001 Ross Avenue, Suite 900, Dallas, TX 75201 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available):

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available):

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

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

(54) Title: BUILDING CUSTOMIZED USER PROFILES BASED ON CONVERSATIONAL DATA

(57) Abstract: In one embodiment, a method includes accessing a plurality of content objects associated with a first user from an on-line social network, accessing a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs, accessing conversational data associated with the first user, determining one or more subjects associated with the first user based on the plurality of content objects and conversational data associated with the first user, and generating a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

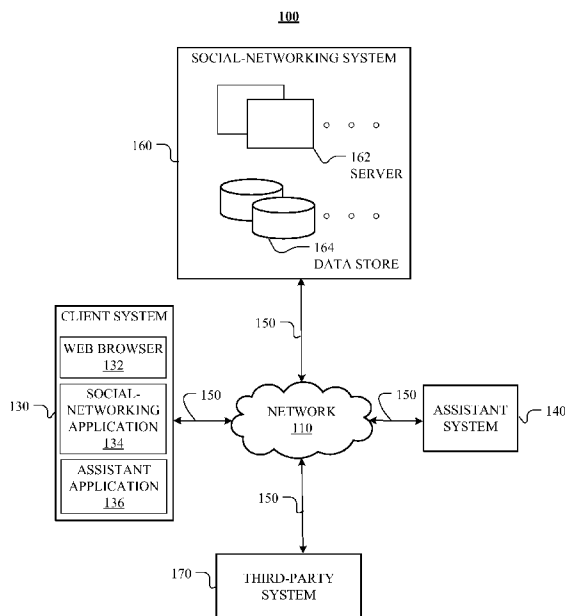


FIG. 1

WO 2019/203867 A1

Published:

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

Building Customized User Profiles Based on Conversational Data

TECHNICAL FIELD

[1] This disclosure generally relates to databases and file management within network environments, and in particular relates to hardware and software for smart assistant systems.

BACKGROUND

[2] An assistant system can provide information or services on behalf of a user based on a combination of user input, location awareness, and the ability to access information from a variety of online sources (such as weather conditions, traffic congestion, news, stock prices, user schedules, retail prices, etc.). The user input may include text (e.g., online chat), especially in an instant messaging application or other applications, voice, images, or a combination of them. The assistant system may perform concierge-type services (e.g., making dinner reservations, purchasing event tickets, making travel arrangements) or provide information based on the user input. The assistant system may also perform management or data-handling tasks based on online information and events without user initiation or interaction. Examples of those tasks that may be performed by an assistant system may include schedule management (e.g., sending an alert to a dinner date that a user is running late due to traffic conditions, update schedules for both parties, and change the restaurant reservation time). The assistant system may be enabled by the combination of computing devices, application programming interfaces (APIs), and the proliferation of applications on user devices.

[3] A social-networking system, which may include a social-networking website, may enable its users (such as persons or organizations) to interact with it and with each other through it. The social-networking system may, with input from a user, create and store in the social-networking system a user profile associated with the user. The user profile may include demographic information, communication-channel information, and information on personal interests of the user. The social-networking system may also, with input from a user, create and store a record of relationships of the user with other users of the social-networking system, as well as provide services (e.g. profile/news feed posts, photo-sharing, event organization, messaging, games, or advertisements) to facilitate social interaction between or among users.

[4] The social-networking system may send over one or more networks content or messages related to its services to a mobile or other computing device of a user. A user may also install software applications on a mobile or other computing device of the user for accessing a user profile of the user and other data within the social-networking system. The social-networking system may generate a personalized set of content objects to display to a user, such as a newsfeed of aggregated stories of other users connected to the user.

SUMMARY OF PARTICULAR EMBODIMENTS

[5] In particular embodiments, the assistant system may assist a user to obtain information or services. The assistant system may enable the user to interact with it with multi-modal user input (such as voice, text, image, video) in stateful and multi-turn conversations to get assistance. The assistant system may create and store a user profile comprising both personal and contextual information associated with the user. In particular embodiments, the assistant system may analyze the user input using natural-language understanding. The analysis may be based on the user profile for more personalized and context-aware understanding. The assistant system may resolve entities associated with the user input based on the analysis. In particular embodiments, the assistant system may interact with different agents to obtain information or services that are associated with the resolved entities. The assistant system may generate a response for the user regarding the information or services by using natural-language generation. Through the interaction with the user, the assistant system may use dialog management techniques to manage and forward the conversation flow with the user. In particular embodiments, the assistant system may further assist the user to effectively and efficiently digest the obtained information by summarizing the information. The assistant system may also assist the user to be more engaging with an online social network by providing tools that help the user interact with the online social network (e.g., creating posts, comments, messages). The assistant system may additionally assist the user to manage different tasks such as keeping track of events. In particular embodiments, the assistant system may proactively execute tasks that are relevant to user interests and preferences based on the user profile without a user input. In particular embodiments, the assistant system may check privacy settings to ensure that accessing a user's profile or other user information and executing different tasks are permitted subject to the user's privacy settings.

[6] In particular embodiments, the assistant system may use a user context engine to build a customized user profile from conversational data (e.g., messages) in combination of other rich semantic information aggregated from various sources. The assistant system may personalize the outputs of other modules of the assistant system based on the customized user profile. In particular embodiments, the other modules of the assistant system may use the customized user profile as the basis of natural-language understanding in response to a user input and execute corresponding tasks based on a more accurately interpreted user input. In particular embodiments, the assistant system may build the customized user profile based on joint processing of one or more of a social graph, a knowledge graph, a concept graph, content objects from an online social network, conversational data which captures contextual information associated with the user, other suitable information related to the user, or any combination thereof. The social graph may include the user's personal information such as location, occupation, connections, etc., whereas the knowledge and concept graphs may capture general knowledge of the world. A knowledge graph may describe real-world entities and their interrelations, define possible classes and relations of entities in a schema, and cover various topical domains. A concept graph may comprise a graph representation for logic based on the semantic networks of artificial intelligence, which may serve as an intermediate language for mapping natural language questions and assertions to a relational database. The content objects from the online social network may be used to identify the user's interests and preferences. The contextual information from the conversational data may capture the particular scenario the user is engaged in, which may provide extra informative cues to understand the user input. In particular embodiments, the user context engine may obtain contextual information by analyzing conversational data in different dialog sessions associated with the user. The customized user profile built by the user context engine may be a tree-structure based representation with different layers characterizing various topics, categories, and entities associated with contextual information. In particular embodiments, the assistant system may use the tree-structure based representation of the user profile to understand the user input in a more personalized and context-aware way with efficiency. Although this disclosure describes building particular user profiles via a particular system in a particular manner, this disclosure contemplates building any suitable user profiles via any suitable system in any suitable manner.

[7] In particular embodiments, the assistant system may access, from an online social network, a plurality of content objects associated with a first user. In particular embodiments, the assistant system may access a baseline profile. The baseline profile may be based on ontology data from one or more information graphs. In particular embodiments, the assistant system may access conversational data associated with the first user (e.g., interactions between the user and an assistant xbot). The assistant system may then determine one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user. In particular embodiments, the assistant system may generate a customized user profile for the first user based on the baseline profile. The customized user profile may comprise one or more confidence scores associated with the respective one or more subjects associated with the first user. The one or more confidence scores may be calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

[8] Certain technical challenges exist for achieving the goal of building customized user profiles based on various data including conversational data. One technical challenge includes incorporating general knowledge of the world into customized user profiles. A solution presented by the embodiments disclosed herein to address the above challenge is learning ontology data from multiple information graphs including knowledge graph, social graph, and concept graph. Another technical challenge includes processing large-scale conversational data efficiently. A solution presented by the embodiments disclosed herein to address this challenge solution is an ensemble-like sliding window analysis. Each sliding window covers N (a pre-specified number) turns at a time for analyzing a conversation (e.g., messages) thread. When $N-1$ turns are completed, a representation for the conversational data processed by the $N-1$ turns is generated. When proceeding to the N -turn sliding window, the embodiments disclosed herein can build upon the representation of the $N-1$ turns for further analysis of the conversational data, thus saving significant processing time.

[9] Certain embodiments disclosed herein may provide one or more technical advantages. A technical advantage of the embodiments may include supporting more personalized and context-aware understanding of users. For example, a user may ask an assistant system to recommend a Chinese restaurant. A more customized user profile with contextual information may

help the assistant system determine at what time with whom at which location (e.g., near home or near work) the user wants to have his/her meal in a Chinese restaurant. Another technical advantage of the embodiments may include integrating information of both head entities and tail entities into customized user profiles. Head entities may be extracted from public sources such as Wikipedia, therefore carrying rich information. By contrast, the information for tail entities may be sparse (e.g., the coffee shop on the corner). Analysis of the conversation data may allow the assistant system to get more information about these tail entities, which are not discussed in public sources (e.g., Wikipedia), but may be discussed considerably in users' conversations or posts on the online social network. Another technical advantage of the embodiments may include easy pruning of insignificant or low-quality entities from a built user profile because of the tree structure. For example, if an entity cannot be put under any categories of the tree-structured representation of the user profile, this entity should be pruned from the user profile. The pruning may be also based on confidence scores of entities. For example, if an entity has a confidence score lower than a threshold score, it may be unimportant to the user and thus be pruned from the user profile. Certain embodiments disclosed herein may provide none, some, or all of the above technical advantages. One or more other technical advantages may be readily apparent to one skilled in the art in view of the figures, descriptions, and claims of the present disclosure.

[10] The embodiments disclosed herein are only examples, and the scope of this disclosure is not limited to them. Particular embodiments may include all, some, or none of the components, elements, features, functions, operations, or steps of the embodiments disclosed above. Embodiments according to the invention are in particular disclosed in the attached claims directed to a method, a storage medium, a system and a computer program product, wherein any feature mentioned in one claim category, e.g. method, can be claimed in another claim category, e.g. system, as well. The dependencies or references back in the attached claims are chosen for formal reasons only. However any subject matter resulting from a deliberate reference back to any previous claims (in particular multiple dependencies) can be claimed as well, so that any combination of claims and the features thereof are disclosed and can be claimed regardless of the dependencies chosen in the attached claims. The subject-matter which can be claimed comprises not only the combinations of features as set out in the attached claims but also any other combination of features in the claims, wherein each feature mentioned in the claims can be

combined with any other feature or combination of other features in the claims. Furthermore, any of the embodiments and features described or depicted herein can be claimed in a separate claim and/or in any combination with any embodiment or feature described or depicted herein or with any of the features of the attached claims.

[11] In an embodiment the method in all embodiments described herein is for use in (or to be executed in) an assistant system for assisting a user to obtain information or services by enabling the user to interact with the assistant system with user input in conversations to get assistance, wherein the user input includes voice, text, image or video or any combination of them. The assistant system is in particular enabled by the combination of computing devices, application programming interfaces (APIs), and the proliferation of applications on user devices.

[12] In an embodiment a method may comprise, by one or more computing systems:

- accessing, from an online social network, a plurality of content objects associated with a first user;
- accessing a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;
- accessing conversational data associated with the first user;
- determining one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and
- generating a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

[13] The plurality of content objects may comprise one or more of:

- news feed posts;
- news feed comments; or
- search history data.

[14] The accessed content objects may be created within a pre-specified time window.

[15] The conversational data may comprise one or more messages created within a pre-specified time window.

[16] Determining the one or more subjects associated with the first user may be based on one or more machine-learning models, wherein the one or more machine-learning models may calculate one or more probabilities of the conversational data being associated with the one or more subjects, respectively.

[17] In an embodiment a method may comprise determining one or more sentiment indications associated with the conversational data, wherein the determining may be based on one or more long-short term memory (LSTM) networks, and wherein the one or more LSTM networks may calculate one or more probabilities of the conversational data being associated with the one or more sentiment indications, respectively.

[18] In an embodiment a method may comprise determining one or more domains associated with the conversational data, wherein the determining may be based on one or more machine-learning models, and wherein the one or more machine-learning models may calculate one or more probabilities of the conversational data being associated with the one or more domains, respectively.

[19] The conversation data may comprise a plurality of messages, and the method may comprise processing the conversational data based on a sliding window analysis over the plurality of messages.

[20] In an embodiment a method may comprise determining contextual information associated with the first user based on the conversational data.

[21] Determining the contextual information associated with the first user may comprise one or more of:

processing meta-data associated with the conversational data based on a machine-learning model, wherein the machine-learning model calculates a probability of the meta-data being associated with the contextual information; or

processing the conversational data based on the machine-learning model, wherein the machine-learning model calculates a probability of the conversational data being associated with the contextual information.

[22] The contextual information may comprise one or more of location information, time stamps, or social connections.

[23] The customized user profile may be based on a tree structure.

[24] The tree structure may comprise one or more layers.

[25] The one or more layers may be organized in a hierarchical structure comprising:

(1) a topic layer comprising one or more topics;

(2) a category layer comprising one or more categories associated with each topic of the topic layer; and

(3) an entity layer comprising one or more entities associated with each category of the category layer.

[26] In an embodiment a method may comprise calculating the one or more confidence scores for the one or more subjects based on a frequency of the respective subject in the plurality of content objects and the conversational data.

[27] Each of the one or more subjects may comprise a topic, a category, or an entity.

[28] In an embodiment a method may comprise determining entity linking information based on the conversational data.

[29] In an embodiment one or more computer-readable non-transitory storage media may embody software that is operable when executed to:

access, from an online social network, a plurality of content objects associated with a first user;

access a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;

access conversational data associated with the first user;

determine one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and

generate a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence

scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

[30] In an embodiment a system may comprise: one or more processors; and a non-transitory memory coupled to the processors comprising instructions executable by the processors, the processors operable when executing the instructions to:

access, from an online social network, a plurality of content objects associated with a first user;

access a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;

access conversational data associated with the first user;

determine one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and

generate a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

[31] Determining the one or more subjects associated with the first user may be based on one or more machine-learning models, wherein the one or more machine-learning models may calculate one or more probabilities of the conversational data being associated with the one or more subjects, respectively, wherein in particular support vector machines (SVM), a regression model, or a deep convolutional neural network (DCNN) may be used.

[32] In an embodiment a method may comprise determining one or more sentiment indications associated with the conversational data, wherein the determining may be based on one or more long-short term memory (LSTM) networks, and wherein the one or more LSTM networks may calculate one or more probabilities of the conversational data being associated with the one or more sentiment indications, respectively.

[33] In an embodiment a method may comprise determining one or more domains associated with the conversational data, wherein the determining may be based on one or more

machine-learning models, and wherein the one or more machine-learning models calculate one or more probabilities of the conversational data being associated with the one or more domains, respectively.

[34] Determining the contextual information associated with the first user may comprise one or more of:

processing meta-data associated with the conversational data based on a machine-learning model, wherein the machine-learning model calculates a probability of the meta-data being associated with the contextual information; or

processing the conversational data based on the machine-learning model, wherein the machine-learning model calculates a probability of the conversational data being associated with the contextual information.

[35] In an embodiment one or more computer-readable non-transitory storage media embodying software that is operable when executed to perform a method according to any of the embodiments herein:

[36] In an embodiment an assistant system for assisting a user to obtain information or services by enabling the user to interact with the assistant system with user input in conversations to get assistance, wherein the user input includes voice, text, image or video or any combination of them, the assistant system in particular being enabled by the combination of computing devices, application programming interfaces (APIs), and the proliferation of applications on user devices, may comprise: one or more processors; and a non-transitory memory coupled to the processors comprising instructions executable by the processors, the processors operable when executing the instructions to perform a method according to any of the embodiments herein.

[37] The assistant system may assist the user by executing at least one or more of the following features or steps:

- create and store a user profile comprising both personal and contextual information associated with the user
- analyze the user input using natural-language understanding, wherein the analysis may be based on the user profile for more personalized and context-aware understanding
- resolve entities associated with the user input based on the analysis

- interact with different agents to obtain information or services that are associated with the resolved entities
- generate a response for the user regarding the information or services by using natural-language generation
- through the interaction with the user, use dialog management techniques to manage and forward the conversation flow with the user
- assist the user to effectively and efficiently digest the obtained information by summarizing the information
- assist the user to be more engaging with an online social network by providing tools that help the user interact with the online social network (e.g., creating posts, comments, messages)
- assist the user to manage different tasks such as keeping track of events
- proactively execute pre-authorized tasks that are relevant to user interests and preferences based on the user profile, at a time relevant for the user, without a user input
- check privacy settings whenever it is necessary to guarantee that accessing user profile and executing different tasks are subject to the user's privacy settings
- automatically identify a user's intent and language register based on a user input, generate a communication content suitable for such intent and language register for the user, detect a change of the user's intent and language register, and dynamically adjust the generation of the communication content to fit the changed intent and language register.

[38] In an embodiment the assistant system comprises at least one or more of the following components:

- a messaging platform to receive a user input based on a text modality from the client system associated with the user and/or to receive user input based on an image or video modality and to process it using optical character recognition techniques within the messaging platform to convert the user input into text,
- an audio speech recognition (ASR) module to receive a user input based on an audio modality (e.g., the user may speak to or send a video including speech) from the client

system associated with the user and to convert the user input based on the audio modality into text,

- an assistant xbot to receive the output of the messaging platform or the ASR module.

[39] In an embodiment the assistant system may comprises a user context engine (USE), wherein the user context engine may store the user profile of the user and wherein the user context engine may access the plurality of content objects associated with the first user from the social-networking system or online social network.

[40] In an embodiment the assistant system may comprises a dialog engine , wherein the dialog engine manages the dialog state and flow of the conversation between the user and the assistant xbot ,

[41] In an embodiment a system may comprise:

at least one client system , in particular an electronic device

at least one assistant system according to any of the embodiments herein,
connected to each other, in particular by a network,

wherein the client system comprises an assistant application for allowing a user at the client system to interact with the assistant system,

wherein the assistant application communicates user input to the assistant system and, based on the user input, the assistant system generates responses and sends the generated responses to the assistant application and the assistant application presents the responses to the user at the client system,

wherein in particular the user input is audio or verbal and the response may be in text or also audio or verbal.

[42] In an embodiment the system may comprise a social-networking system,

wherein the client system may comprise in particular a social-networking application for accessing the social networking system.

BRIEF DESCRIPTION OF THE DRAWINGS

[43] FIG. 1 illustrates an example network environment associated with an assistant system.

[44] FIG. 2 illustrates an example architecture of the assistant system.

[45] FIG. 3 illustrates an example diagram flow of responding to a user request by the assistant system.

[46] FIG. 4 illustrates an example architecture of building a customized user profile.

[47] FIG. 5 illustrates an example tree-structured user profile.

[48] FIG. 6 illustrates an example method for building a customized user profile.

[49] FIG. 7 illustrates an example social graph.

[50] FIG. 8 illustrates an example view of an embedding space.

[51] FIG. 9 illustrates an example artificial neural network.

[52] FIG. 10 illustrates an example computer system.

DESCRIPTION OF EXAMPLE EMBODIMENTS

System Overview

[53] **FIG. 1** illustrates an example network environment 100 associated with an assistant system. Network environment 100 includes a client system 130, an assistant system 140, a social-networking system 160, and a third-party system 170 connected to each other by a network 110. Although FIG. 1 illustrates a particular arrangement of a client system 130, an assistant system 140, a social-networking system 160, a third-party system 170, and a network 110, this disclosure contemplates any suitable arrangement of a client system 130, an assistant system 140, a social-networking system 160, a third-party system 170, and a network 110. As an example and not by way of limitation, two or more of a client system 130, a social-networking system 160, an assistant system 140, and a third-party system 170 may be connected to each other directly, bypassing a network 110. As another example, two or more of a client system 130, an assistant system 140, a social-networking system 160, and a third-party system 170 may be physically or logically co-located with each other in whole or in part. Moreover, although FIG. 1 illustrates a particular number of client systems 130, assistant systems 140, social-networking systems 160, third-party systems 170, and networks 110, this disclosure contemplates any suitable number of client systems 130, assistant systems 140, social-networking systems 160, third-party systems 170, and networks 110. As an example and not by way of limitation, network environment 100 may include multiple client systems 130, assistant systems 140, social-networking systems 160, third-party systems 170, and networks 110.

[54] This disclosure contemplates any suitable network 110. As an example and not by way of limitation, one or more portions of a network 110 may include an ad hoc network, an intranet, an extranet, a virtual private network (VPN), a local area network (LAN), a wireless LAN (WLAN), a wide area network (WAN), a wireless WAN (WWAN), a metropolitan area network (MAN), a portion of the Internet, a portion of the Public Switched Telephone Network (PSTN), a cellular telephone network, or a combination of two or more of these. A network 110 may include one or more networks 110.

[55] Links 150 may connect a client system 130, an assistant system 140, a social-networking system 160, and a third-party system 170 to a communication network 110 or to each other. This disclosure contemplates any suitable links 150. In particular embodiments, one or more links 150 include one or more wireline (such as for example Digital Subscriber Line (DSL) or Data Over Cable Service Interface Specification (DOCSIS)), wireless (such as for example Wi-Fi or Worldwide Interoperability for Microwave Access (WiMAX)), or optical (such as for example Synchronous Optical Network (SONET) or Synchronous Digital Hierarchy (SDH)) links. In particular embodiments, one or more links 150 each include an ad hoc network, an intranet, an extranet, a VPN, a LAN, a WLAN, a WAN, a WWAN, a MAN, a portion of the Internet, a portion of the PSTN, a cellular technology-based network, a satellite communications technology-based network, another link 150, or a combination of two or more such links 150. Links 150 need not necessarily be the same throughout a network environment 100. One or more first links 150 may differ in one or more respects from one or more second links 150.

[56] In particular embodiments, a client system 130 may be an electronic device including hardware, software, or embedded logic components or a combination of two or more such components and capable of carrying out the appropriate functionalities implemented or supported by a client system 130. As an example and not by way of limitation, a client system 130 may include a computer system such as a desktop computer, notebook or laptop computer, netbook, a tablet computer, e-book reader, GPS device, camera, personal digital assistant (PDA), handheld electronic device, cellular telephone, smartphone, smart speaker, other suitable electronic device, or any suitable combination thereof. In particular embodiments, a client system 130 may be a smart assistant device. More information on smart assistant devices may be found in U.S. Patent Application No. 15/949011, filed 09 April 2018, U.S. Patent Application No.

62/655751, filed 10 April 2018, U.S. Patent Application No. 29/631910, filed 03 January 2018, U.S. Patent Application No. 29/631747, filed 02 January 2018, U.S. Patent Application No. 29/631913, filed 03 January 2018, and U.S. Patent Application No. 29/631914, filed 03 January 2018, which are incorporated by reference. This disclosure contemplates any suitable client systems 130. A client system 130 may enable a network user at a client system 130 to access a network 110. A client system 130 may enable its user to communicate with other users at other client systems 130.

[57] In particular embodiments, a client system 130 may include a web browser 132, such as MICROSOFT INTERNET EXPLORER, GOOGLE CHROME or MOZILLA FIREFOX, and may have one or more add-ons, plug-ins, or other extensions, such as TOOLBAR or YAHOO TOOLBAR. A user at a client system 130 may enter a Uniform Resource Locator (URL) or other address directing a web browser 132 to a particular server (such as server 162, or a server associated with a third-party system 170), and the web browser 132 may generate a Hyper Text Transfer Protocol (HTTP) request and communicate the HTTP request to server. The server may accept the HTTP request and communicate to a client system 130 one or more Hyper Text Markup Language (HTML) files responsive to the HTTP request. The client system 130 may render a web interface (e.g. a webpage) based on the HTML files from the server for presentation to the user. This disclosure contemplates any suitable source files. As an example and not by way of limitation, a web interface may be rendered from HTML files, Extensible Hyper Text Markup Language (XHTML) files, or Extensible Markup Language (XML) files, according to particular needs. Such interfaces may also execute scripts such as, for example and without limitation, those written in JAVASCRIPT, JAVA, MICROSOFT SILVERLIGHT, combinations of markup language and scripts such as AJAX (Asynchronous JAVASCRIPT and XML), and the like. Herein, reference to a web interface encompasses one or more corresponding source files (which a browser may use to render the web interface) and vice versa, where appropriate.

[58] In particular embodiments, a client system 130 may include a social-networking application 134 installed on the client system 130. A user at a client system 130 may use the social-networking application 134 to access an online social network. The user at the client system 130 may use the social-networking application 134 to communicate with the user's social connections (e.g., friends, followers, followed accounts, contacts, etc.). The user at the client system 130 may

also use the social-networking application 134 to interact with a plurality of content objects (e.g., posts, news articles, ephemeral content, etc.) on the online social network. As an example and not by way of limitation, the user may browse trending topics and breaking news using the social-networking application 134.

[59] In particular embodiments, a client system 130 may include an assistant application 136. A user at a client system 130 may use the assistant application 136 to interact with the assistant system 140. In particular embodiments, the assistant application 136 may comprise a stand-alone application. In particular embodiments, the assistant application 136 may be integrated into the social-networking application 134 or another suitable application (e.g., a messaging application). In particular embodiments, the assistant application 136 may be also integrated into the client system 130, an assistant hardware device, or any other suitable hardware devices. In particular embodiments, the assistant application 136 may be accessed via the web browser 132. In particular embodiments, the user may provide input via different modalities. As an example and not by way of limitation, the modalities may include audio, text, image, video, etc. The assistant application 136 may communicate the user input to the assistant system 140. Based on the user input, the assistant system 140 may generate responses. The assistant system 140 may send the generated responses to the assistant application 136. The assistant application 136 may then present the responses to the user at the client system 130. The presented responses may be based on different modalities such as audio, text, image, and video. As an example and not by way of limitation, the user may verbally ask the assistant application 136 about the traffic information (i.e., via an audio modality). The assistant application 136 may then communicate the request to the assistant system 140. The assistant system 140 may accordingly generate the result and send it back to the assistant application 136. The assistant application 136 may further present the result to the user in text.

[60] In particular embodiments, an assistant system 140 may assist users to retrieve information from different sources. The assistant system 140 may also assist user to request services from different service providers. In particular embodiments, the assist system 140 may receive a user request for information or services via the assistant application 136 in the client system 130. The assist system 140 may use natural-language understanding to analyze the user request based on user profile and other relevant information. The result of the analysis may comprise different entities associated with an online social network. The assistant system 140 may

then retrieve information or request services associated with these entities. In particular embodiments, the assistant system 140 may interact with the social-networking system 160 and/or third-party system 170 when retrieving information or requesting services for the user. In particular embodiments, the assistant system 140 may generate a personalized communication content for the user using natural-language generating techniques. The personalized communication content may comprise, for example, the retrieved information or the status of the requested services. In particular embodiments, the assistant system 140 may enable the user to interact with it regarding the information or services in a stateful and multi-turn conversation by using dialog-management techniques. The functionality of the assistant system 140 is described in more detail in the discussion of FIG. 2 below.

[61] In particular embodiments, the social-networking system 160 may be a network-addressable computing system that can host an online social network. The social-networking system 160 may generate, store, receive, and send social-networking data, such as, for example, user-profile data, concept-profile data, social-graph information, or other suitable data related to the online social network. The social-networking system 160 may be accessed by the other components of network environment 100 either directly or via a network 110. As an example and not by way of limitation, a client system 130 may access the social-networking system 160 using a web browser 132, or a native application associated with the social-networking system 160 (e.g., a mobile social-networking application, a messaging application, another suitable application, or any combination thereof) either directly or via a network 110. In particular embodiments, the social-networking system 160 may include one or more servers 162. Each server 162 may be a unitary server or a distributed server spanning multiple computers or multiple datacenters. Servers 162 may be of various types, such as, for example and without limitation, web server, news server, mail server, message server, advertising server, file server, application server, exchange server, database server, proxy server, another server suitable for performing functions or processes described herein, or any combination thereof. In particular embodiments, each server 162 may include hardware, software, or embedded logic components or a combination of two or more such components for carrying out the appropriate functionalities implemented or supported by server 162. In particular embodiments, the social-networking system 160 may include one or more data stores 164. Data stores 164 may be used to store various types of information. In particular

embodiments, the information stored in data stores 164 may be organized according to specific data structures. In particular embodiments, each data store 164 may be a relational, columnar, correlation, or other suitable database. Although this disclosure describes or illustrates particular types of databases, this disclosure contemplates any suitable types of databases. Particular embodiments may provide interfaces that enable a client system 130, a social-networking system 160, or a third-party system 170 to manage, retrieve, modify, add, or delete, the information stored in data store 164.

[62] In particular embodiments, the social-networking system 160 may store one or more social graphs in one or more data stores 164. In particular embodiments, a social graph may include multiple nodes—which may include multiple user nodes (each corresponding to a particular user) or multiple concept nodes (each corresponding to a particular concept)—and multiple edges connecting the nodes. The social-networking system 160 may provide users of the online social network the ability to communicate and interact with other users. In particular embodiments, users may join the online social network via the social-networking system 160 and then add connections (e.g., relationships) to a number of other users of the social-networking system 160 whom they want to be connected to. Herein, the term “friend” may refer to any other user of the social-networking system 160 with whom a user has formed a connection, association, or relationship via the social-networking system 160.

[63] In particular embodiments, the social-networking system 160 may provide users with the ability to take actions on various types of items or objects, supported by the social-networking system 160. As an example and not by way of limitation, the items and objects may include groups or social networks to which users of the social-networking system 160 may belong, events or calendar entries in which a user might be interested, computer-based applications that a user may use, transactions that allow users to buy or sell items via the service, interactions with advertisements that a user may perform, or other suitable items or objects. A user may interact with anything that is capable of being represented in the social-networking system 160 or by an external system of a third-party system 170, which is separate from the social-networking system 160 and coupled to the social-networking system 160 via a network 110.

[64] In particular embodiments, the social-networking system 160 may be capable of linking a variety of entities. As an example and not by way of limitation, the social-networking

system 160 may enable users to interact with each other as well as receive content from third-party systems 170 or other entities, or to allow users to interact with these entities through an application programming interfaces (API) or other communication channels.

[65] In particular embodiments, a third-party system 170 may include one or more types of servers, one or more data stores, one or more interfaces, including but not limited to APIs, one or more web services, one or more content sources, one or more networks, or any other suitable components, e.g., that servers may communicate with. A third-party system 170 may be operated by a different entity from an entity operating the social-networking system 160. In particular embodiments, however, the social-networking system 160 and third-party systems 170 may operate in conjunction with each other to provide social-networking services to users of the social-networking system 160 or third-party systems 170. In this sense, the social-networking system 160 may provide a platform, or backbone, which other systems, such as third-party systems 170, may use to provide social-networking services and functionality to users across the Internet.

[66] In particular embodiments, a third-party system 170 may include a third-party content object provider. A third-party content object provider may include one or more sources of content objects, which may be communicated to a client system 130. As an example and not by way of limitation, content objects may include information regarding things or activities of interest to the user, such as, for example, movie show times, movie reviews, restaurant reviews, restaurant menus, product information and reviews, or other suitable information. As another example and not by way of limitation, content objects may include incentive content objects, such as coupons, discount tickets, gift certificates, or other suitable incentive objects.

[67] In particular embodiments, the social-networking system 160 also includes user-generated content objects, which may enhance a user's interactions with the social-networking system 160. User-generated content may include anything a user can add, upload, send, or "post" to the social-networking system 160. As an example and not by way of limitation, a user communicates posts to the social-networking system 160 from a client system 130. Posts may include data such as status updates or other textual data, location information, photos, videos, links, music or other similar data or media. Content may also be added to the social-networking system 160 by a third-party through a "communication channel," such as a newsfeed or stream.

[68] In particular embodiments, the social-networking system 160 may include a variety of servers, sub-systems, programs, modules, logs, and data stores. In particular embodiments, the social-networking system 160 may include one or more of the following: a web server, action logger, API-request server, relevance-and-ranking engine, content-object classifier, notification controller, action log, third-party-content-object-exposure log, inference module, authorization/privacy server, search module, advertisement-targeting module, user-interface module, user-profile store, connection store, third-party content store, or location store. The social-networking system 160 may also include suitable components such as network interfaces, security mechanisms, load balancers, failover servers, management-and-network-operations consoles, other suitable components, or any suitable combination thereof. In particular embodiments, the social-networking system 160 may include one or more user-profile stores for storing user profiles. A user profile may include, for example, biographic information, demographic information, behavioral information, social information, or other types of descriptive information, such as work experience, educational history, hobbies or preferences, interests, affinities, or location. Interest information may include interests related to one or more categories. Categories may be general or specific. As an example and not by way of limitation, if a user “likes” an article about a brand of shoes the category may be the brand, or the general category of “shoes” or “clothing.” A connection store may be used for storing connection information about users. The connection information may indicate users who have similar or common work experience, group memberships, hobbies, educational history, or are in any way related or share common attributes. The connection information may also include user-defined connections between different users and content (both internal and external). A web server may be used for linking the social-networking system 160 to one or more client systems 130 or one or more third-party systems 170 via a network 110. The web server may include a mail server or other messaging functionality for receiving and routing messages between the social-networking system 160 and one or more client systems 130. An API-request server may allow a third-party system 170 to access information from the social-networking system 160 by calling one or more APIs. An action logger may be used to receive communications from a web server about a user’s actions on or off the social-networking system 160. In conjunction with the action log, a third-party-content-object log may be maintained of user exposures to third-party-content objects. A notification controller may provide information

regarding content objects to a client system 130. Information may be pushed to a client system 130 as notifications, or information may be pulled from a client system 130 responsive to a request received from a client system 130. Authorization servers may be used to enforce one or more privacy settings of the users of the social-networking system 160. A privacy setting of a user determines how particular information associated with a user can be shared. The authorization server may allow users to opt in to or opt out of having their actions logged by the social-networking system 160 or shared with other systems (e.g., a third-party system 170), such as, for example, by setting appropriate privacy settings. Third-party-content-object stores may be used to store content objects received from third parties, such as a third-party system 170. Location stores may be used for storing location information received from client systems 130 associated with users. Advertisement-pricing modules may combine social information, the current time, location information, or other suitable information to provide relevant advertisements, in the form of notifications, to a user.

Assistant Systems

[69] FIG. 2 illustrates an example architecture of the assistant system 140. In particular embodiments, the assistant system 140 may assist a user to obtain information or services. The assistant system 140 may enable the user to interact with it with multi-modal user input (such as voice, text, image, video) in stateful and multi-turn conversations to get assistance. The assistant system 140 may create and store a user profile comprising both personal and contextual information associated with the user. In particular embodiments, the assistant system 140 may analyze the user input using natural-language understanding. The analysis may be based on the user profile for more personalized and context-aware understanding. The assistant system 140 may resolve entities associated with the user input based on the analysis. In particular embodiments, the assistant system 140 may interact with different agents to obtain information or services that are associated with the resolved entities. The assistant system 140 may generate a response for the user regarding the information or services by using natural-language generation. Through the interaction with the user, the assistant system 140 may use dialog management techniques to manage and forward the conversation flow with the user. In particular embodiments, the assistant system 140 may further assist the user to effectively and efficiently digest the obtained information by summarizing the information. The assistant system 140 may also assist the user to be more

engaging with an online social network by providing tools that help the user interact with the online social network (e.g., creating posts, comments, messages). The assistant system 140 may additionally assist the user to manage different tasks such as keeping track of events. In particular embodiments, the assistant system 140 may proactively execute pre-authorized tasks that are relevant to user interests and preferences based on the user profile, at a time relevant for the user, without a user input. In particular embodiments, the assistant system 140 may check privacy settings whenever it is necessary to guarantee that accessing user profile and executing different tasks are subject to the user's privacy settings.

[70] In particular embodiments, the assistant system 140 may receive a user input from the assistant application 136 in the client system 130 associated with the user. If the user input is based on a text modality, the assistant system 140 may receive it at a messaging platform 205. If the user input is based on an audio modality (e.g., the user may speak to the assistant application 136 or send a video including speech to the assistant application 136), the assistant system 140 may process it using an audio speech recognition (ASR) module 210 to convert the user input into text. If the user input is based on an image or video modality, the assistant system 140 may process it using optical character recognition techniques within the messaging platform 205 to convert the user input into text. The output of the messaging platform 205 or the ASR module 210 may be received at an assistant xbot 215.

[71] In particular embodiments, the assistant xbot 215 may be a type of chat bot. The assistant xbot 215 may comprise a programmable service channel, which may be a software code, logic, or routine that functions as a personal assistant to the user. The assistant xbot 215 may work as the user's portal to the assistant system 140. The assistant xbot 215 may therefore be considered as a type of conversational agent. In particular embodiments, the assistant xbot 215 may send the textual user input to a natural-language understanding (NLU) module 220 to interpret the user input. In particular embodiments, the NLU module 220 may get information from a user context engine 225 and a semantic information aggregator 230 to accurately understand the user input. The user context engine 225 may store the user profile of the user. The user profile of the user may comprise user-profile data including demographic information, social information, and contextual information associated with the user. The user-profile data may also include user interests and preferences on a plurality of topics, aggregated through conversations on news feed, search logs,

messaging platform 205, etc. The usage of user profile may be protected behind a privacy check module 245 to ensure that a user's information can be used only for his/her benefit, and not shared with anyone else. The semantic information aggregator 230 may provide ontology data associated with a plurality of predefined domains, intents, and slots to the NLU module 220. In particular embodiments, a domain may denote a social context of interaction, e.g., education. An intent may indicate a purpose of a user interacting with the assistant system 140. A slot may represent a basic semantic entity. For example, a slot for "pizza" may be dish. The semantic information aggregator 230 may additionally extract information from a social graph, a knowledge graph, and a concept graph, and retrieve user profile from the user context engine 225. The semantic information aggregator 230 may further process information from these different sources by determining what information to aggregate, annotating n-grams of the user input, ranking the n-grams with confidence scores based on the aggregated information, formulating the ranked n-grams into features that can be used by the NLU module 220 for understanding the user input. Based on the output of the user context engine 225 and the semantic information aggregator 230, the NLU module 220 may identify a domain, an intent, and one or more slots from the user input in a personalized and context-aware manner. As an example and not by way of limitation, a user input may comprise "show me how to get to the Starbucks". The NLU module 220 may identify the particular Starbucks that the user wants to go based on the user's personal information and the associated contextual information. In particular embodiments, the NLU module 220 may comprise a lexicon of language and a parser and grammar rules to partition sentences into an internal representation. The NLU module 220 may also comprise one or more programs that perform naive semantics or stochastic semantic analysis to the use of pragmatics to understand a user input. In particular embodiments, the parser may be based on a deep learning architecture comprising multiple long-short term memory (LSTM) networks. As an example and not by way of limitation, the parser may be based on a recurrent neural network grammar (RNNG) model, which is a type of recurrent and recursive LSTM algorithm.

[72] In particular embodiments, the identified domain, intent, and one or more slots from the NLU module 220 may be sent to a dialog engine 235. In particular embodiments, the dialog engine 235 may manage the dialog state and flow of the conversation between the user and the assistant xbot 215. The dialog engine 235 may additionally store previous conversations between

the user and the assistant xbot 215. In particular embodiments, the dialog engine 235 may communicate with an entity resolution module 240 to resolve entities associated with the one or more slots, which supports the dialog engine 235 to forward the flow of the conversation between the user and the assistant xbot 215. In particular embodiments, the entity resolution module 240 may access the social graph, the knowledge graph, and the concept graph when resolving the entities. Entities may include, for example, unique users or concepts, each of which may have a unique identifier (ID). As an example and not by way of limitation, the knowledge graph may comprise a plurality of entities. Each entity may comprise a single record associated with one or more attribute values. The particular record may be associated with a unique entity identifier. Each record may have diverse values for an attribute of the entity. Each attribute value may be associated with a confidence probability. A confidence probability for an attribute value represents a probability that the value is accurate for the given attribute. Each attribute value may be also associated with a semantic weight. A semantic weight for an attribute value may represent how the value semantically appropriate for the given attribute considering all the available information. For example, the knowledge graph may comprise an entity of a movie “The Martian” (2015), which includes information that has been extracted from multiple content sources (e.g., Facebook, Wikipedia, CSM, IMDb, Metacritics, Rovi, TMDb, and Rotten Tomatoes), and then deduped, resolved, and fused to generate the single unique record for the knowledge graph. The entity may be associated with a space attribute value which indicates the genre of the movie “The Martian” (2015). The entity resolution module 240 may additionally request user profile of the user associated with the user input from the user context engine 225. In particular embodiments, the entity resolution module 240 may communicate with a privacy check module 245 to guarantee that the resolving of the entities does not violate privacy policies. In particular embodiments, the privacy check module 245 may use an authorization/privacy server to enforce privacy policies. As an example and not by way of limitation, an entity to be resolved may be another user who specifies in his/her privacy settings that his/her identity should not be searchable on the online social network, and thus the entity resolution module 240 may not return that user’s identifier in response to a request. Based on the information obtained from the social graph, knowledge graph, concept graph, and user profile, and subject to applicable privacy policies, the entity resolution module 240 may therefore accurately resolve the entities associated with the user input in a personalized and

context-aware manner. In particular embodiments, each of the resolved entities may be associated with one or more identifiers hosted by the social-networking system 160. As an example and not by way of limitation, an identifier may comprise a unique user identifier (ID). In particular embodiments, each of the resolved entities may be also associated with a confidence score.

[73] In particular embodiments, the dialog engine 235 may communicate with different agents based on the identified intent and domain, and the resolved entities. In particular embodiments, the agents may comprise first-party agents 250 and third-party agents 255. In particular embodiments, first-party agents 250 may comprise internal agents that are accessible and controllable by the assistant system 140 (e.g. agents associated with services provided by the online social network (Messenger, Instagram)). In particular embodiments, third-party agents 255 may comprise external agents that the assistant system 140 has no control over (e.g., music streams agents (Spotify), ticket sales agents (Ticketmaster)). The first-party agents 250 may be associated with first-party providers 260 that provide content objects and/or services hosted by the social-networking system 160. The third-party agents 255 may be associated with third-party providers 265 that provide content objects and/or services hosted by the third-party system 170.

[74] In particular embodiments, the communication from the dialog engine 235 to the first-party agents 250 may comprise requesting particular content objects and/or services provided by the first-party providers 260. As a result, the first-party agents 250 may retrieve the requested content objects from the first-party providers 260 and/or execute tasks that command the first-party providers 260 to perform the requested services. In particular embodiments, the communication from the dialog engine 235 to the third-party agents 255 may comprise requesting particular content objects and/or services provided by the third-party providers 265. As a result, the third-party agents 255 may retrieve the requested content objects from the third-party providers 265 and/or execute tasks that command the third-party providers 265 to perform the requested services. The third-party agents 255 may access the privacy check module 245 to guarantee no privacy violations before interacting with the third-party providers 265. As an example and not by way of limitation, the user associated with the user input may specify in his/her privacy settings that his/her profile information is invisible to any third-party content providers. Therefore, when retrieving content objects associated with the user input from the third-party providers 265, the

third-party agents 255 may complete the retrieval without revealing to the third-party providers 265 which user is requesting the content objects.

[75] In particular embodiments, each of the first-party agents 250 or third-party agents 255 may be designated for a particular domain. As an example and not by way of limitation, the domain may comprise weather, transportation, music, etc. In particular embodiments, the assistant system 140 may use a plurality of agents collaboratively to respond to a user input. As an example and not by way of limitation, the user input may comprise “direct me to my next meeting.” The assistant system 140 may use a calendar agent to retrieve the location of the next meeting. The assistant system 140 may then use a navigation agent to direct the user to the next meeting.

[76] In particular embodiments, each of the first-party agents 250 or third-party agents 255 may retrieve user profile from the user context engine 225 to execute tasks in a personalized and context-aware manner. As an example and not by way of limitation, a user input may comprise “book me a ride to the airport.” A transportation agent may execute the task of booking the ride. The transportation agent may retrieve user profile of the user from the user context engine 225 before booking the ride. For example, the user profile may indicate that the user prefers taxis, so the transportation agent may book a taxi for the user. As another example, the contextual information associated with the user profile may indicate that the user is in a hurry so the transportation agent may book a ride from a ride-sharing service (e.g., Uber, Lyft) for the user since it may be faster to get a car from a ride-sharing service than a taxi company. In particular embodiment, each of the first-party agents 250 or third-party agents 255 may take into account other factors when executing tasks. As an example and not by way of limitation, other factors may comprise price, rating, efficiency, partnerships with the online social network, etc.

[77] In particular embodiments, the dialog engine 235 may communicate with a conversational understanding composer (CU composer) 270. The dialog engine 235 may send the requested content objects and/or the statuses of the requested services to the CU composer 270. In particular embodiments, the dialog engine 235 may send the requested content objects and/or the statuses of the requested services as a $\langle k, c, u, d \rangle$ tuple, in which k indicates a knowledge source, c indicates a communicative goal, u indicates a user model, and d indicates a discourse model. In particular embodiments, the CU composer 270 may comprise a natural-language generator (NLG) 271 and a user interface (UI) payload generator 272. The natural-language

generator 271 may generate a communication content based on the output of the dialog engine 235. In particular embodiments, the NLG 271 may comprise a content determination component, a sentence planner, and a surface realization component. The content determination component may determine the communication content based on the knowledge source, communicative goal, and the user's expectations. As an example and not by way of limitation, the determining may be based on a description logic. The description logic may comprise, for example, three fundamental notions which are individuals (representing objects in the domain), concepts (describing sets of individuals), and roles (representing binary relations between individuals or concepts). The description logic may be characterized by a set of constructors that allow the natural-language generator 271 to build complex concepts/roles from atomic ones. In particular embodiments, the content determination component may perform the following tasks to determine the communication content. The first task may comprise a translation task, in which the input to the natural-language generator 271 may be translated to concepts. The second task may comprise a selection task, in which relevant concepts may be selected among those resulted from the translation task based on the user model. The third task may comprise a verification task, in which the coherence of the selected concepts may be verified. The fourth task may comprise an instantiation task, in which the verified concepts may be instantiated as an executable file that can be processed by the natural-language generator 271. The sentence planner may determine the organization of the communication content to make it human understandable. The surface realization component may determine specific words to use, the sequence of the sentences, and the style of the communication content. The UI payload generator 272 may determine a preferred modality of the communication content to be presented to the user. In particular embodiments, the CU composer 270 may communicate with the privacy check module 245 to make sure the generation of the communication content follows the privacy policies. In particular embodiments, the CU composer 270 may retrieve user profile from the user context engine 225 when generating the communication content and determining the modality of the communication content. As a result, the communication content may be more natural, personalized, and context-aware for the user. As an example and not by way of limitation, the user profile may indicate that the user likes short sentences in conversations so the generated communication content may be based on short sentences. As another example and not by way of limitation, the contextual information associated

with the user profile may indicate that the user is using a device that only outputs audio signals so the UI payload generator 272 may determine the modality of the communication content as audio.

[78] In particular embodiments, the CU composer 270 may send the generated communication content to the assistant xbot 215. In particular embodiments, the assistant xbot 215 may send the communication content to the messaging platform 205. The messaging platform 205 may further send the communication content to the client system 130 via the assistant application 136. In alternative embodiments, the assistant xbot 215 may send the communication content to a text-to-speech (TTS) module 275. The TTS module 275 may convert the communication content to an audio clip. The TTS module 275 may further send the audio clip to the client system 130 via the assistant application 136.

[79] In particular embodiments, the assistant xbot 215 may interact with a proactive inference layer 280 without receiving a user input. The proactive inference layer 280 may infer user interests and preferences based on user profile that is retrieved from the user context engine 225. In particular embodiments, the proactive inference layer 280 may further communicate with proactive agents 285 regarding the inference. The proactive agents 285 may execute proactive tasks based on the inference. As an example and not by way of limitation, the proactive tasks may comprise sending content objects or providing services to the user. In particular embodiments, each proactive task may be associated with an agenda item. The agenda item may comprise a recurring item such as a daily digest. The agenda item may also comprise a one-time item. In particular embodiments, a proactive agent 285 may retrieve user profile from the user context engine 225 when executing the proactive task. Therefore, the proactive agent 285 may execute the proactive task in a personalized and context-aware manner. As an example and not by way of limitation, the proactive inference layer may infer that the user likes the band Maroon 5 and the proactive agent 285 may generate a recommendation of Maroon 5's new song/album to the user.

[80] In particular embodiments, the proactive agent 285 may generate candidate entities associated with the proactive task based on user profile. The generation may be based on a straightforward backend query using deterministic filters to retrieve the candidate entities from a structured data store. The generation may be alternatively based on a machine-learning model that is trained based on user profile, entity attributes, and relevance between users and entities. As an

example and not by way of limitation, the machine-learning model may be based on support vector machines (SVM). As another example and not by way of limitation, the machine-learning model may be based on a regression model. As another example and not by way of limitation, the machine-learning model may be based on a deep convolutional neural network (DCNN). In particular embodiments, the proactive agent 285 may also rank the generated candidate entities based on user profile and the content associated with the candidate entities. The ranking may be based on the similarities between a user's interests and the candidate entities. As an example and not by way of limitation, the assistant system 140 may generate a feature vector representing a user's interest and feature vectors representing the candidate entities. The assistant system 140 may then calculate similarity scores (e.g., based on cosine similarity) between the feature vector representing the user's interest and the feature vectors representing the candidate entities. The ranking may be alternatively based on a ranking model that is trained based on user feedback data.

[81] In particular embodiments, the proactive task may comprise recommending the candidate entities to a user. The proactive agent 285 may schedule the recommendation, thereby associating a recommendation time with the recommended candidate entities. The recommended candidate entities may be also associated with a priority and an expiration time. In particular embodiments, the recommended candidate entities may be sent to a proactive scheduler. The proactive scheduler may determine an actual time to send the recommended candidate entities to the user based on the priority associated with the task and other relevant factors (e.g., clicks and impressions of the recommended candidate entities). In particular embodiments, the proactive scheduler may then send the recommended candidate entities with the determined actual time to an asynchronous tier. The asynchronous tier may temporarily store the recommended candidate entities as a job. In particular embodiments, the asynchronous tier may send the job to the dialog engine 235 at the determined actual time for execution. In alternative embodiments, the asynchronous tier may execute the job by sending it to other surfaces (e.g., other notification services associated with the social-networking system 160). In particular embodiments, the dialog engine 235 may identify the dialog intent, state, and history associated with the user. Based on the dialog intent, the dialog engine 235 may select some candidate entities among the recommended candidate entities to send to the client system 130. In particular embodiments, the dialog state and history may indicate if the user is engaged in an ongoing conversation with the assistant xbot 215.

If the user is engaged in an ongoing conversation and the priority of the task of recommendation is low, the dialog engine 235 may communicate with the proactive scheduler to reschedule a time to send the selected candidate entities to the client system 130. If the user is engaged in an ongoing conversation and the priority of the task of recommendation is high, the dialog engine 235 may initiate a new dialog session with the user in which the selected candidate entities may be presented. As a result, the interruption of the ongoing conversation may be prevented. When it is determined that sending the selected candidate entities is not interruptive to the user, the dialog engine 235 may send the selected candidate entities to the CU composer 270 to generate a personalized and context-aware communication content comprising the selected candidate entities, subject to the user's privacy settings. In particular embodiments, the CU composer 270 may send the communication content to the assistant xbot 215 which may then send it to the client system 130 via the messaging platform 205 or the TTS module 275.

[82] In particular embodiments, the assistant xbot 215 may communicate with a proactive agent 285 in response to a user input. As an example and not by way of limitation, the user may ask the assistant xbot 215 to set up a reminder. The assistant xbot 215 may request a proactive agent 285 to set up such reminder and the proactive agent 285 may proactively execute the task of reminding the user at a later time.

[83] In particular embodiments, the assistant system 140 may comprise a summarizer 290. The summarizer 290 may provide customized news feed summaries to a user. In particular embodiments, the summarizer 290 may comprise a plurality of meta agents. The plurality of meta agents may use the first-party agents 250, third-party agents 255, or proactive agents 285 to generate news feed summaries. In particular embodiments, the summarizer 290 may retrieve user interests and preferences from the proactive inference layer 280. The summarizer 290 may then retrieve entities associated with the user interests and preferences from the entity resolution module 240. The summarizer 290 may further retrieve user profile from the user context engine 225. Based on the information from the proactive inference layer 280, the entity resolution module 240, and the user context engine 225, the summarizer 290 may generate personalized and context-aware summaries for the user. In particular embodiments, the summarizer 290 may send the summaries to the CU composer 270. The CU composer 270 may process the summaries and send the

processing results to the assistant xbot 215. The assistant xbot 215 may then send the processed summaries to the client system 130 via the messaging platform 205 or the TTS module 275.

[84] FIG. 3 illustrates an example diagram flow of responding to a user request by the assistant system 140. In particular embodiments, the assistant xbot 215 may access a request manager 305 upon receiving the user request. The request manager 305 may comprise a context extractor 306 and a conversational understanding object generator (CU object generator) 307. The context extractor 306 may extract contextual information associated with the user request. The CU object generator 307 may generate particular content objects relevant to the user request. In particular embodiments, the request manager 305 may store the contextual information and the generated content objects in data store 310 which is a particular data store implemented in the assistant system 140.

[85] In particular embodiments, the request manager 305 may send the generated content objects to the NLU module 220. The NLU module 220 may perform a plurality of steps to process the content objects. At step 221, the NLU module 220 may generate a whitelist for the content objects. At step 222, the NLU module 220 may perform a featurization based on the whitelist. At step 223, the NLU module 220 may perform domain classification/selection based on the features resulted from the featurization. The domain classification/selection results may be further processed based on two related procedures. At step 224a, the NLU module 220 may process the domain classification/selection result using an intent classifier. The intent classifier may determine the user's intent associated with the user request. As an example and not by way of limitation, the intent classifier may be based on a machine-learning model that may take the domain classification/selection result as input and calculate a probability of the input being associated with a particular predefined intent. At step 224b, the NLU module may process the domain classification/selection result using a meta-intent classifier. The meta-intent classifier may determine categories that describe the user's intent. As an example and not by way of limitation, the meta-intent classifier may be based on a machine-learning model that may take the domain classification/selection result as input and calculate a probability of the input being associated with a particular predefined meta-intent. At step 225a, the NLU module 220 may use a slot tagger to annotate one or more slots associated with the user request. At step 225b, the NLU module 220 may use a meta slot tagger to annotate one or more slots for the classification result from the meta-

intent classifier. As an example and not by way of limitation, a user request may comprise “change 500 dollars in my account to Japanese yen.” The intent classifier may take the user request as input and formulate it into a vector. The intent classifier may then calculate probabilities of the user request being associated with different predefined intents based on a vector comparison between the vector representing the user request and the vectors representing different predefined intents. In a similar manner, the slot tagger may take the user request as input and formulate each word into a vector. The intent classifier may then calculate probabilities of each word being associated with different predefined slots based on a vector comparison between the vector representing the word and the vectors representing different predefined slots. The intent of the user may be classified as “changing money”. The slots of the user request may comprise “500”, “dollars”, “account”, and “Japanese yen”. The meta-intent of the user may be classified as “financial service”. The meta slot may comprise “finance”.

[86] In particular embodiments, the NLU module 220 may improve the domain classification/selection of the content objects by extracting semantic information from the semantic information aggregator 230. In particular embodiments, the semantic information aggregator 230 may aggregate semantic information in the following way. The semantic information aggregator 230 may first retrieve information from the user context engine 225. In particular embodiments, the user context engine 225 may comprise offline aggregators 226 and an online inference service 227. The offline aggregators 226 may process a plurality of data associated with the user that are collected from a prior time window. As an example and not by way of limitation, the data may include news feed posts/comments, interactions with news feed posts/comments, Instagram posts/comments, search history, etc. that are collected from a prior 90-day window. The processing result may be stored in the user context engine 225 as part of the user profile. The online inference service 227 may analyze the conversational data associated with the user that are received by the assistant system 140 at a current time. The analysis result may be stored in the user context engine 225 also as part of the user profile. In particular embodiments, the semantic information aggregator 230 may then process the retrieved information, i.e., user profile, from the user context engine 225 in the following steps. At step 231, the semantic information aggregator 230 may process the retrieved information from the user context engine 225 based on natural-language processing (NLP). At step 232, the processing result may be annotated with entities by an entity tagger. Based

on the annotations, the semantic information aggregator 230 may generate dictionaries for the retrieved information at step 233. At step 234, the semantic information aggregator 230 may rank the entities tagged by the entity tagger. In particular embodiments, the semantic information aggregator 230 may communicate with different graphs 330 including social graph, knowledge graph, and concept graph to extract ontology data that is relevant to the retrieved information from the user context engine 225. In particular embodiments, the semantic information aggregator 230 may aggregate user profile, the ranked entities, and the information from the graphs 330. The semantic information aggregator 230 may then send the aggregated information to the NLU module 220 to facilitate the domain classification/selection.

[87] In particular embodiments, the output of the NLU module 220 may be sent to a co-reference module 315 to interpret references of the content objects associated with the user request. The co-reference module 315 may comprise reference creation 316 and reference resolution 317. In particular embodiments, the reference creation 316 may create references for entities determined by the NLU module 220. The reference resolution 317 may resolve these references accurately. As an example and not by way of limitation, a user request may comprise “find me the nearest Walmart and direct me there”. The co-reference module 315 may interpret “there” as “the nearest Walmart”. In particular embodiments, the co-reference module 315 may access the user context engine 225 and the dialog engine 235 when necessary to interpret references with improved accuracy.

[88] In particular embodiments, the identified domains, intents, meta-intents, slots, and meta slots, along with the resolved references may be sent to the entity resolution module 240 to resolve relevant entities. In particular embodiments, the entity resolution module 240 may comprise domain entity resolution 241 and generic entity resolution 242. The domain entity resolution 241 may resolve the entities by categorizing the slots and meta slots into different domains. In particular embodiments, entities may be resolved based on the ontology data extracted from the graphs 330. The ontology data may comprise the structural relationship between different slots/meta-slots and domains. The ontology may also comprise information of how the slots/meta-slots may be grouped, related within a hierarchy where the higher level comprises the domain, and subdivided according to similarities and differences. The generic entity resolution 242 may resolve the entities by categorizing the slots and meta slots into different generic topics. In particular

embodiments, the resolving may be also based on the ontology data extracted from the graphs 330. The ontology data may comprise the structural relationship between different slots/meta-slots and generic topics. The ontology may also comprise information of how the slots/meta-slots may be grouped, related within a hierarchy where the higher level comprises the topic, and subdivided according to similarities and differences. As an example and not by way of limitation, in response to the input of an inquiry of the advantages of a Tesla car, the generic entity resolution 242 may resolve a Tesla car as vehicle and the domain entity resolution 241 may resolve the Tesla car as electric car.

[89] In particular embodiments, the output of the entity resolution module 240 may be sent to the dialog engine 235 to forward the flow of the conversation with the user. The dialog engine 235 may comprise dialog intent resolution 236 and dialog state update/ranker 237. In particular embodiments, the dialog intent resolution 236 may resolve the user intent associated with the current dialog session. In particular embodiments, the dialog state update/ranker 237 may update/rank the state of the current dialog session. As an example and not by way of limitation, the dialog state update/ranker 237 may update the dialog state as “completed” if the dialog session is over. As another example and not by way of limitation, the dialog state update/ranker 237 may rank the dialog state based on a priority associated with it.

[90] In particular embodiments, the dialog engine 235 may communicate with a task completion module 335 about the dialog intent and associated content objects. The task completion module 335 may comprise an action selection component 336. In particular embodiments, the dialog engine 235 may additionally check against dialog policies 320 regarding the dialog state. The dialog policies 320 may comprise generic policy 321 and domain specific policies 322, both of which may guide how to select the next system action based on the dialog state. In particular embodiments, the task completion module 335 may communicate with dialog policies 320 to obtain the guidance of the next system action. In particular embodiments, the action selection component 336 may therefore select an action based on the dialog intent, the associated content objects, and the guidance from dialog policies 320.

[91] In particular embodiments, the output of the task completion module 335 may be sent to the CU composer 270. In alternative embodiments, the selected action may require one or more agents 340 to be involved. As a result, the task completion module 335 may inform the agents

340 about the selected action. Meanwhile, the dialog engine 235 may receive an instruction to update the dialog state. As an example and not by way of limitation, the update may comprise awaiting agents' response. In particular embodiments, the CU composer 270 may generate a communication content for the user using the NLG 271 based on the output of the task completion module 335. The CU composer 270 may also determine a modality of the generated communication content using the UI payload generator 272. Since the generated communication content may be considered as a response to the user request, the CU composer 270 may additionally rank the generated communication content using a response ranker 273. As an example and not by way of limitation, the ranking may indicate the priority of the response.

[92] In particular embodiments, the output of the CU composer 270 may be sent to a response manager 325. The response manager 325 may perform different tasks including storing/updating the dialog state 326 retrieved from data store 310 and generating responses 327. In particular embodiments, the generated response and the communication content may be sent to the assistant xbot 215. In alternative embodiments, the output of the CU composer 270 may be additionally sent to the TTS module 275 if the determined modality of the communication content is audio. The speech generated by the TTS module 275 and the response generated by the response manager 325 may be then sent to the assistant xbot 215.

Building Customized User Profiles Based on Conversational Data

[93] In particular embodiments, an assistant system 140 may use a user context engine 225 to build a customized user profile from conversational data (e.g., messages) in combination of other rich semantic information aggregated from various sources. The assistant system 140 may personalize the outputs of other modules of the assistant system 140 based on the customized user profile. In particular embodiments, the other modules of the assistant system 140 may use the customized user profile as the basis of natural-language understanding in response to a user input and execute corresponding tasks based on a more accurately interpreted user input. In particular embodiments, the assistant system 140 may build the customized user profile based on joint processing of one or more of a social graph, a knowledge graph, a concept graph, content objects from an online social network, conversational data which captures contextual information associated with the user, other suitable information related to the user, or any combination thereof. The social graph may include the user's personal information such as location, occupation,

connections, etc., whereas the knowledge and concept graphs may capture general knowledge of the world. The knowledge graph may describe real-world entities and their interrelations, define possible classes and relations of entities in a schema, and cover various topical domains. The concept graph may comprise a graph representation for logic based on the semantic networks of artificial intelligence, which may serve as an intermediate language for mapping natural language questions and assertions to a relational database. The content objects from the online social network may be used to identify the user's interests and preferences. The contextual information from the conversational data may capture the particular scenario the user is engaged in, which may provide extra informative cues to understand the user input. In particular embodiments, the user context engine 225 may obtain contextual information by analyzing conversational data in different dialog sessions associated with the user. The customized user profile built by the user context engine 225 may be a tree-structure based representation with different layers characterizing various topics, categories, and entities associated with contextual information. In particular embodiments, the assistant system 140 may use the tree-structure based representation of the user profile to understand the user input in a more personalized and context-aware way with efficiency. Although this disclosure describes building particular user profiles via a particular system in a particular manner, this disclosure contemplates building any suitable user profiles via any suitable system in any suitable manner.

[94] In particular embodiments, the assistant system 140 may access, from an online social network, a plurality of content objects associated with a first user. In particular embodiments, the assistant system 140 may access a baseline profile. The baseline profile may be based on ontology data from one or more information graphs. In particular embodiments, the assistant system 140 may access conversational data associated with the first user (e.g., interactions between the user and an assistant xbot 215). The assistant system 140 may then determine one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user. In particular embodiments, the assistant system 140 may generate a customized user profile for the first user based on the baseline profile. The customized user profile may comprise one or more confidence scores associated with the respective one or more subjects associated with the first user. The one

or more confidence scores may be calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

[95] FIG. 4 illustrates an example architecture of building a customized user profile. In particular embodiments, the user context engine 225 may access the plurality of content objects 401 from the social-networking system 160. The accessed content objects 401 may be created within a pre-specified time window. As an example and not by way of limitation, the pre-specified time window may be one week, one month, three months, or any other suitable time window that can be specified by the user context engine 225. The social-networking system 160 may store the plurality of content objects 401 in the one or more data stores 164. In particular embodiments, the plurality of content objects 401 may comprise one or more of news feed posts, news feed comments, search history data, other suitable content, or any combination thereof. In particular embodiments, the user context engine 225 may additionally extract ontology data 402 from one or more information graphs 330. Extracting ontology data 402 from multiple information graphs including a knowledge graph, social graph, and concept graph may be an effective solution for addressing the technical challenge of incorporating general knowledge of the world into user profiles. In particular embodiments, the user context engine 225 may obtain conversational data 403 from the dialog engine 235 and/or the data store 310 of the assistant system 140 illustrated in FIG. 3. The conversational data 403 may comprise one or more messages created within a pre-specified time window. As an example and not by way of limitation, the pre-specified time window may be ten minutes, one hour, one day, one week, one month, or any other suitable time window that can be specified by the user context engine 225. In particular embodiments, the conversational data 403 from the dialog engine 235 may comprise messages between the first user and the assistant xbot 215. The usage of these messages is subject to the first user's privacy settings. In particular embodiments, the conversational data 403 from the data store 310 of the assistant system 140 may comprise messages between the first user and one or more second users. The usage of these messages is subject to both the first user's and second users' privacy settings. Using conversational data 403 may result in a technical advantage of integrating information of both head entities and tail entities into user profiles. Head entities may be extracted from public sources such as Wikipedia, therefore carrying rich information. By contrast, the information for tail entities may be sparse (e.g., the coffee shop on the corner). Analysis of the conversation data 403 may allow

the assistant system 140 to get more information about these tail entities, which are not discussed in public sources (e.g., Wikipedia), but may be discussed considerably in users' conversations. Based on an analysis of the content objects 401, ontology data 402, and conversational data 403, the user context engine 225 may build a user profile that includes both personal information and contextual information for the first user. Although this disclosure describes building particular user profiles using particular modules in a particular manner, this disclosure contemplates building any suitable user profiles using any suitable modules in any suitable manner.

[96] In particular embodiments, the assistant system 140 may determine the one or more subjects associated with the first user based on one or more machine-learning models. In particular embodiments, each of the one or more subjects may comprise a topic, a category, or an entity. As an example and not by way of limitation, the one or more machine-learning model may be based on one or more of support vector machines (SVM), a regression model, or a deep convolutional neural network (DCNN). The machine-learning models may take the conversational data 403 as input and calculate a probability of the conversational data 403 being associated with a particular predefined topic, category, or entity. In particular embodiments, the assistant system 140 may calculate the one or more confidence scores for the one or more subjects based on a frequency of the respective subject in the plurality of content objects 401 and the conversational data 403. As an example and not by way of limitation, within a week time window, a user may have written ten posts and may have mentioned BMW car in three of these ten posts. In addition, the user may have ten conversations with other users and may have mentioned BMW car in seven of these conversations. Accordingly, the confidence score may be calculated as $\frac{3+7}{10+10} = 0.5$. Although this disclosure describes determining particular subjects and their confidence scores in a particular manner, this disclosure contemplates determining any suitable subjects and their confidence scores in any suitable manner.

[97] In particular embodiments, the assistant system 140 may perform different analysis on the conversational data 403. The conversation data 403 may comprise a plurality of messages, and therefore the assistant system 104 may process the conversational data 403 based on a sliding window analysis over the plurality of messages. A technical challenge for processing the conversational data 403 may include efficient processing. The solution presented by the embodiments disclosed herein is an ensemble-like sliding window analysis. Each sliding window

covers N (a pre-specified number) turns at a time for analyzing a conversation (e.g., messages) thread. When $N-1$ turns are completed, a representation for the conversational data 403 processed by the $N-1$ turns is generated. When proceeding to the N -turn sliding window, the embodiments disclosed herein can simply build upon the representation of the $N-1$ turns for further analysis of the conversational data 403, thus saving significant processing time. In particular embodiments, the assistant system 140 may determine one or more sentiment indications associated with the conversational data 403. In particular embodiments, the assistant system 140 may perform a different analysis on the conversational data 403 based on a LSTM-based model comprising a unidirectional LSTM layer, a dense layer, and a softmax layer. The input of the LSTM-based model may comprise the conversational data 403 and the output may be probabilities of the conversational data 403 being associated with different sentiment indications. A sentiment indication may comprise the attitude or emotional reaction of a user characterized in the conversational data 403. As an example and not by way of limitation, sentiment indications may comprise emotional states such as “angry”, “sad”, and “happy”. In particular embodiments, the assistant system 140 may determine one or more domains associated with the conversational data 403. As an example and not by way of limitation, a domain may comprise education, entertainment, or politics. In particular embodiments, the assistant system 140 may determine the one or more domains based on one or more machine-learning models. As an example and not by way of limitation, the one or more machine-learning models may be based on one or more of support vector machines (SVM), a regression model, or a deep convolutional neural network (DCNN). The machine-learning models may take the conversational data 403 as input and calculate a probability of the conversational data 403 being associated with a particular predefined domain. In particular embodiments, the assistant system 140 may determine entity linking information based on the conversational data 403. As an example and not by way of limitation, the entity linking information may comprise information indicating a particular text is explicitly or implicitly linked to a particular ID of an entity. For example, a user may explicitly tag an entity in a post/comment, which leads to explicit entity linking information. As another example, the user may mention the name of the entity without tagging the entity in the post/comment, which leads to implicit entity linking information. In this case, the assistant system 140 may still be able to infer that the entity should be linked to that post/comment along with a confidence score for the entity linking information. More information

on entity linking may be found in U.S. Patent Application No. 15/355500, filed 18 November 2016, and U.S. Patent Application No. 15/827622, filed 30 November 2017, which are incorporated by reference. Although this disclosure describes performing particular analysis on particular data in a particular manner, this disclosure contemplates performing any suitable analysis on any suitable data in any suitable manner.

[98] In particular embodiments, the assistant system 140 may determine contextual information associated with the first user based on the conversational data 403. As an example and not by way of limitation, the contextual information may comprise one or more of location information, time stamps, social connections, other suitable information, or any combination thereof. In particular embodiments, the assistant system 140 may determine the contextual information associated with the first user by processing meta-data associated with the conversational data 403 based on a machine-learning model. The meta-data may describe the conversational data 403 and comprise information about the conversational data. As an example and not by way of limitation, a user may have a conversation with another user as “I’m thinking about watching a movie together tomorrow night for our second date. What time works for you?”. The meta-data associated with the conversation may comprise “date night”, “romance”, “entertainment”, and “planning”. In particular embodiments, the assistant system 140 may alternatively determine the contextual information associated with the first user by processing the conversational data 403 based on the machine-learning model. In particular embodiments, the assistant system 140 may also determine the contextual information associated with the first user by processing both the meta-data associated with the conversational data 403 and the conversational data 403 based on the machine-learning model. In particular embodiments, the machine-learning model may be based on a LSTM-based model comprising a unidirectional LSTM layer, a dense layer, and a softmax layer. The input of the LSTM-based model may comprise the conversational data 403 and the output may be probabilities of the conversational data 403 being associated with different types of contextual information. A technical advantage of the embodiments disclosed herein may include supporting more personalized and context-aware understanding of users. For example, a user may ask the assistant system 140 to recommend a Chinese restaurant. The built user profile with contextual information may help the assistant system 140 determine at what time with whom at which location (e.g., near home or near work)

the user wants to have his/her meal in a Chinese restaurant. Although this disclosure describes determining particular contextual information based on particular data in a particular manner, this disclosure contemplates determining any suitable contextual information based on any suitable data in any suitable manner.

[99] FIG. 5 illustrates an example tree-structured user profile. In particular embodiments, the customized user profile may be based on a tree structure. The tree structure may comprise one or more layers. In particular embodiments, the one or more layers may be organized in a hierarchical structure. The hierarchical structure may comprise: (1) a topic layer comprising one or more topics; (2) a category layer comprising one or more categories associated with each topic of the topic layer; and (3) an entity layer comprising one or more entities associated with each category of the category layer. As an example and not by way of limitation, the user profile displayed in FIG. 5 may comprise a topic layer 510, a category layer 520, and an entity layer 530. For example, two topics of the topic layer 510 are displayed in FIG. 5, which are “movie” 511 and “travel” 512. Each topic is associated with a confidence score respectively, i.e., 0.88 and 0.52, which may indicate that “movie” 511 is more significantly associated with the user than “travel” 512. Under the topic of “movie” 511 there are two categories in the category layer 520 including “action” 521 and “sci-fi” 522, each being associated with a confidence score. Under the category “action” 521, there are two entities in the entity layer 530 including “Die Hard” 531 and “Kill Bill” 532, which are names of two movies. “Die Hard” 531 and “Kill Bill” 532 are associated with confidence scores, respectively. Similar structure can be observed for “sci-fi” 522 associated with “Star Wars” 533 and “Avatar” 534, which are names of another two sci-fi movies. As for another topic “travel” 512, similar structure indicates that it is associated with “city” 523 and “landscape” 524; that “city” 523 is associated with “Milan” 535 and “Paris” 536; and that “landscape” 524 is associated with “Yellowstone” 537 which is a national park of the USA. In particular embodiments, each entity may be additionally associated with contextual information. As displayed in FIG. 5, each entity is associated with contextual information 540. For example, the contextual information 540 for entity “Star Wars” 533 comprises “AMC Metreon, 02/17/18, with friend Dan”. “AMC Metreon” is a movie theater in San Francisco, which indicates location information. “02/17/18” indicates the time stamp, and “with friend Dan” indicates the social connection. As the user profile is a tree structure, it may be used in different ways. The topic layer

510 and the category layer 520 may be used in a heuristic approach in identifying a user's general interests and preferences. The entity layer 530 may be processed by the user context engine 225 based on a sparse convolutional neural network, which may result in a personalized user-profile representation. The tree-structured user profile may result a technical advantage of easy pruning of insignificant entities from a built user profile. For example, if an entity cannot be put under any categories of the tree-structured user profile, this entity should be pruned from the user profile. The pruning may be also based on confidence scores of entities. For example, if an entity has a confidence score lower than a threshold score, it may be unimportant to the user and thus be pruned from the user profile. Although this disclosure describes particular user profiles with particular structures in a particular manner, this disclosure contemplates any suitable user profiles with any suitable structures in any suitable manner.

[100] In particular embodiments, the customized user profile may be used by different modules of the assistant system 140. As an example and not by way of limitation, the proactive agents 285 may use a heuristic approach to identify user interests based on the customized user profile. The proactive agents 285 may also take the customized user profile as input and train a personalized machine-learning model. The personalized machine-learning model, in combination with other data, may be used to optimize the previously built customized user profile, thus resulting in an even better customized user profile. As another example and not by way of limitation, the proactive agents 285 may provide better services to the user. For example, the customized user profile may indicate that under the context of being at work, the user always listens to classical music. Hence, the proactive agents 285 may suggest classical music to the user when the user is at work. As another example and not by way of limitation, the NLU module 220 may adjust the confidence scores associated with its output based on the customized user profile. For example, a user may ask the assistant system 140 "how can I get to Lucca?". Based on the baseline profile, the NLU module 220 may determine a <city> slot for "Lucca" with a confidence score 0.9 since it is a city in Italy. However, the customized user profile may indicate that the user is significantly associated with a restaurant named Lucca in San Francisco (e.g., the user liked the webpage of the restaurant several times). Therefore, the NLU module 220 may adjust the confidence score of the <city> slot for "Lucca" to 0.2 and determine a <restaurant> slot for "Lucca" with a confidence score of 0.78. As another example and not by way of limitation, the CU composer 270 may

generate personalized communication content based on the customized user profile. For example, a user's brother Aaron posted a video on his news feed. Based on the baseline profile, the CU composer 270 may generate a communication content for the user as "Aaron posted a video". However, the customized user profile of the user may indicate that the user always chats with his brother John using a nickname (e.g. "Boy"). As a result, the CU composer 270 may modify the generation of the communication content to be more personalized as "Boy posted a video". Although this disclosure describes particular use cases of the user profile by particular modules in a particular manner, this disclosure contemplates any suitable use cases of the user profile by any suitable modules in any suitable manner.

[101] FIG. 6 illustrates an example method 600 for building a customized user profile. The method may begin at step 610, where the assistant system 140 may access, from an online social network, a plurality of content objects 401 associated with a first user. At step 620, the assistant system 140 may access a baseline profile, wherein the baseline profile is based on ontology data 402 from one or more information graphs 330. At step 630, the assistant system 140 may access conversational data 403 associated with the first user. At step 640, the assistant system 140 may determine one or more subjects associated with the first user based on the plurality of content objects 401 associated with the first user and the conversational data 403 associated with the first user. At step 650, the assistant system 140 may generate a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects 401 associated with the first user and the conversational data 403 associated with the first user. Particular embodiments may repeat one or more steps of the method of FIG. 6, where appropriate. Although this disclosure describes and illustrates particular steps of the method of FIG. 6 as occurring in a particular order, this disclosure contemplates any suitable steps of the method of FIG. 6 occurring in any suitable order. Moreover, although this disclosure describes and illustrates an example method for building a user profile, including the particular steps of the method of FIG. 6, this disclosure contemplates any suitable method for building a user profile, including any suitable steps, which may include all, some, or none of the steps of the method of FIG. 6, where appropriate. Furthermore, although this disclosure describes and illustrates particular components,

devices, or systems carrying out particular steps of the method of FIG. 6, this disclosure contemplates any suitable combination of any suitable components, devices, or systems carrying out any suitable steps of the method of FIG. 6.

Social Graphs

[102] FIG. 7 illustrates an example social graph 700. In particular embodiments, the social-networking system 160 may store one or more social graphs 700 in one or more data stores. In particular embodiments, the social graph 700 may include multiple nodes—which may include multiple user nodes 702 or multiple concept nodes 704—and multiple edges 706 connecting the nodes. Each node may be associated with a unique entity (i.e., user or concept), each of which may have a unique identifier (ID), such as a unique number or username. The example social graph 700 illustrated in FIG. 7 is shown, for didactic purposes, in a two-dimensional visual map representation. In particular embodiments, a social-networking system 160, a client system 130, an assistant system 170, or a third-party system 170 may access the social graph 700 and related social-graph information for suitable applications. The nodes and edges of the social graph 700 may be stored as data objects, for example, in a data store (such as a social-graph database). Such a data store may include one or more searchable or queryable indexes of nodes or edges of the social graph 700.

[103] In particular embodiments, a user node 702 may correspond to a user of the social-networking system 160 or the assistant system 170. As an example and not by way of limitation, a user may be an individual (human user), an entity (e.g., an enterprise, business, or third-party application), or a group (e.g., of individuals or entities) that interacts or communicates with or over the social-networking system 160 or the assistant system 170. In particular embodiments, when a user registers for an account with the social-networking system 160, the social-networking system 160 may create a user node 702 corresponding to the user, and store the user node 702 in one or more data stores. Users and user nodes 702 described herein may, where appropriate, refer to registered users and user nodes 702 associated with registered users. In addition or as an alternative, users and user nodes 702 described herein may, where appropriate, refer to users that have not registered with the social-networking system 160. In particular embodiments, a user node 702 may be associated with information provided by a user or information gathered by various systems, including the social-networking system 160. As an example and not by way of limitation,

a user may provide his or her name, profile picture, contact information, birth date, sex, marital status, family status, employment, education background, preferences, interests, or other demographic information. In particular embodiments, a user node 702 may be associated with one or more data objects corresponding to information associated with a user. In particular embodiments, a user node 702 may correspond to one or more web interfaces.

[104] In particular embodiments, a concept node 704 may correspond to a concept. As an example and not by way of limitation, a concept may correspond to a place (such as, for example, a movie theater, restaurant, landmark, or city); a website (such as, for example, a website associated with the social-networking system 160 or a third-party website associated with a web-application server); an entity (such as, for example, a person, business, group, sports team, or celebrity); a resource (such as, for example, an audio file, video file, digital photo, text file, structured document, or application) which may be located within the social-networking system 160 or on an external server, such as a web-application server; real or intellectual property (such as, for example, a sculpture, painting, movie, game, song, idea, photograph, or written work); a game; an activity; an idea or theory; another suitable concept; or two or more such concepts. A concept node 704 may be associated with information of a concept provided by a user or information gathered by various systems, including the social-networking system 160 and the assistant system 170. As an example and not by way of limitation, information of a concept may include a name or a title; one or more images (e.g., an image of the cover page of a book); a location (e.g., an address or a geographical location); a website (which may be associated with a URL); contact information (e.g., a phone number or an email address); other suitable concept information; or any suitable combination of such information. In particular embodiments, a concept node 704 may be associated with one or more data objects corresponding to information associated with concept node 704. In particular embodiments, a concept node 704 may correspond to one or more web interfaces.

[105] In particular embodiments, a node in the social graph 700 may represent or be represented by a web interface (which may be referred to as a “profile interface”). Profile interfaces may be hosted by or accessible to the social-networking system 160 or the assistant system 170. Profile interfaces may also be hosted on third-party websites associated with a third-party system 170. As an example and not by way of limitation, a profile interface corresponding to a particular

external web interface may be the particular external web interface and the profile interface may correspond to a particular concept node 704. Profile interfaces may be viewable by all or a selected subset of other users. As an example and not by way of limitation, a user node 702 may have a corresponding user-profile interface in which the corresponding user may add content, make declarations, or otherwise express himself or herself. As another example and not by way of limitation, a concept node 704 may have a corresponding concept-profile interface in which one or more users may add content, make declarations, or express themselves, particularly in relation to the concept corresponding to concept node 704.

[106] In particular embodiments, a concept node 704 may represent a third-party web interface or resource hosted by a third-party system 170. The third-party web interface or resource may include, among other elements, content, a selectable or other icon, or other inter-actable object (which may be implemented, for example, in JavaScript, AJAX, or PHP codes) representing an action or activity. As an example and not by way of limitation, a third-party web interface may include a selectable icon such as “like,” “check-in,” “eat,” “recommend,” or another suitable action or activity. A user viewing the third-party web interface may perform an action by selecting one of the icons (e.g., “check-in”), causing a client system 130 to send to the social-networking system 160 a message indicating the user’s action. In response to the message, the social-networking system 160 may create an edge (e.g., a check-in-type edge) between a user node 702 corresponding to the user and a concept node 704 corresponding to the third-party web interface or resource and store edge 706 in one or more data stores.

[107] In particular embodiments, a pair of nodes in the social graph 700 may be connected to each other by one or more edges 706. An edge 706 connecting a pair of nodes may represent a relationship between the pair of nodes. In particular embodiments, an edge 706 may include or represent one or more data objects or attributes corresponding to the relationship between a pair of nodes. As an example and not by way of limitation, a first user may indicate that a second user is a “friend” of the first user. In response to this indication, the social-networking system 160 may send a “friend request” to the second user. If the second user confirms the “friend request,” the social-networking system 160 may create an edge 706 connecting the first user’s user node 702 to the second user’s user node 702 in the social graph 700 and store edge 706 as social-graph information in one or more of data stores 167. In the example of FIG. 7, the social graph 700

includes an edge 706 indicating a friend relation between user nodes 702 of user “A” and user “B” and an edge indicating a friend relation between user nodes 702 of user “C” and user “B.” Although this disclosure describes or illustrates particular edges 706 with particular attributes connecting particular user nodes 702, this disclosure contemplates any suitable edges 706 with any suitable attributes connecting user nodes 702. As an example and not by way of limitation, an edge 706 may represent a friendship, family relationship, business or employment relationship, fan relationship (including, e.g., liking, etc.), follower relationship, visitor relationship (including, e.g., accessing, viewing, checking-in, sharing, etc.), subscriber relationship, superior/subordinate relationship, reciprocal relationship, non-reciprocal relationship, another suitable type of relationship, or two or more such relationships. Moreover, although this disclosure generally describes nodes as being connected, this disclosure also describes users or concepts as being connected. Herein, references to users or concepts being connected may, where appropriate, refer to the nodes corresponding to those users or concepts being connected in the social graph 700 by one or more edges 706.

[108] In particular embodiments, an edge 706 between a user node 702 and a concept node 704 may represent a particular action or activity performed by a user associated with user node 702 toward a concept associated with a concept node 704. As an example and not by way of limitation, as illustrated in FIG. 7, a user may “like,” “attended,” “played,” “listened,” “cooked,” “worked at,” or “watched” a concept, each of which may correspond to an edge type or subtype. A concept-profile interface corresponding to a concept node 704 may include, for example, a selectable “check in” icon (such as, for example, a clickable “check in” icon) or a selectable “add to favorites” icon. Similarly, after a user clicks these icons, the social-networking system 160 may create a “favorite” edge or a “check in” edge in response to a user’s action corresponding to a respective action. As another example and not by way of limitation, a user (user “C”) may listen to a particular song (“Imagine”) using a particular application (SPOTIFY, which is an online music application). In this case, the social-networking system 160 may create a “listened” edge 706 and a “used” edge (as illustrated in FIG. 7) between user nodes 702 corresponding to the user and concept nodes 704 corresponding to the song and application to indicate that the user listened to the song and used the application. Moreover, the social-networking system 160 may create a “played” edge 706 (as illustrated in FIG. 7) between concept nodes 704 corresponding to the song

and the application to indicate that the particular song was played by the particular application. In this case, “played” edge 706 corresponds to an action performed by an external application (SPOTIFY) on an external audio file (the song “Imagine”). Although this disclosure describes particular edges 706 with particular attributes connecting user nodes 702 and concept nodes 704, this disclosure contemplates any suitable edges 706 with any suitable attributes connecting user nodes 702 and concept nodes 704. Moreover, although this disclosure describes edges between a user node 702 and a concept node 704 representing a single relationship, this disclosure contemplates edges between a user node 702 and a concept node 704 representing one or more relationships. As an example and not by way of limitation, an edge 706 may represent both that a user likes and has used at a particular concept. Alternatively, another edge 706 may represent each type of relationship (or multiples of a single relationship) between a user node 702 and a concept node 704 (as illustrated in FIG. 7 between user node 702 for user “E” and concept node 704 for “SPOTIFY”).

[109] In particular embodiments, the social-networking system 160 may create an edge 706 between a user node 702 and a concept node 704 in the social graph 700. As an example and not by way of limitation, a user viewing a concept-profile interface (such as, for example, by using a web browser or a special-purpose application hosted by the user’s client system 130) may indicate that he or she likes the concept represented by the concept node 704 by clicking or selecting a “Like” icon, which may cause the user’s client system 130 to send to the social-networking system 160 a message indicating the user’s liking of the concept associated with the concept-profile interface. In response to the message, the social-networking system 160 may create an edge 706 between user node 702 associated with the user and concept node 704, as illustrated by “like” edge 706 between the user and concept node 704. In particular embodiments, the social-networking system 160 may store an edge 706 in one or more data stores. In particular embodiments, an edge 706 may be automatically formed by the social-networking system 160 in response to a particular user action. As an example and not by way of limitation, if a first user uploads a picture, watches a movie, or listens to a song, an edge 706 may be formed between user node 702 corresponding to the first user and concept nodes 704 corresponding to those concepts. Although this disclosure describes forming particular edges 706 in particular manners, this disclosure contemplates forming any suitable edges 706 in any suitable manner.

Vector Spaces and Embeddings

[110] FIG. 8 illustrates an example view of a vector space 800. In particular embodiments, an object or an n-gram may be represented in a d -dimensional vector space, where d denotes any suitable number of dimensions. Although the vector space 800 is illustrated as a three-dimensional space, this is for illustrative purposes only, as the vector space 800 may be of any suitable dimension. In particular embodiments, an n-gram may be represented in the vector space 800 as a vector referred to as a term embedding. Each vector may comprise coordinates corresponding to a particular point in the vector space 800 (i.e., the terminal point of the vector). As an example and not by way of limitation, vectors 810, 820, and 880 may be represented as points in the vector space 800, as illustrated in FIG. 8. An n-gram may be mapped to a respective vector representation. As an example and not by way of limitation, n-grams t_1 and t_2 may be mapped to vectors \vec{v}_1 and \vec{v}_2 in the vector space 800, respectively, by applying a function $\vec{\pi}$ defined by a dictionary, such that $\vec{v}_1 = \vec{\pi}(t_1)$ and $\vec{v}_2 = \vec{\pi}(t_2)$. As another example and not by way of limitation, a dictionary trained to map text to a vector representation may be utilized, or such a dictionary may be itself generated via training. As another example and not by way of limitation, a model, such as Word2vec, may be used to map an n-gram to a vector representation in the vector space 800. In particular embodiments, an n-gram may be mapped to a vector representation in the vector space 800 by using a machine learning model (e.g., a neural network). The machine learning model may have been trained using a sequence of training data (e.g., a corpus of objects each comprising n-grams).

[111] In particular embodiments, an object may be represented in the vector space 800 as a vector referred to as a feature vector or an object embedding. As an example and not by way of limitation, objects e_1 and e_2 may be mapped to vectors \vec{v}_1 and \vec{v}_2 in the vector space 800, respectively, by applying a function $\vec{\pi}$, such that $\vec{v}_1 = \vec{\pi}(e_1)$ and $\vec{v}_2 = \vec{\pi}(e_2)$. In particular embodiments, an object may be mapped to a vector based on one or more properties, attributes, or features of the object, relationships of the object with other objects, or any other suitable information associated with the object. As an example and not by way of limitation, a function $\vec{\pi}$ may map objects to vectors by feature extraction, which may start from an initial set of measured data and build derived values (e.g., features). As an example and not by way of limitation, an object comprising a video or an image may be mapped to a vector by using an algorithm to detect or

isolate various desired portions or shapes of the object. Features used to calculate the vector may be based on information obtained from edge detection, corner detection, blob detection, ridge detection, scale-invariant feature transformation, edge direction, changing intensity, autocorrelation, motion detection, optical flow, thresholding, blob extraction, template matching, Hough transformation (e.g., lines, circles, ellipses, arbitrary shapes), or any other suitable information. As another example and not by way of limitation, an object comprising audio data may be mapped to a vector based on features such as a spectral slope, a tonality coefficient, an audio spectrum centroid, an audio spectrum envelope, a Mel-frequency cepstrum, or any other suitable information. In particular embodiments, when an object has data that is either too large to be efficiently processed or comprises redundant data, a function $\vec{\pi}$ may map the object to a vector using a transformed reduced set of features (e.g., feature selection). In particular embodiments, a function $\vec{\pi}$ may map an object e to a vector $\vec{\pi}(e)$ based on one or more n-grams associated with object e . Although this disclosure describes representing an n-gram or an object in a vector space in a particular manner, this disclosure contemplates representing an n-gram or an object in a vector space in any suitable manner.

[112] In particular embodiments, the social-networking system 160 may calculate a similarity metric of vectors in vector space 800. A similarity metric may be a cosine similarity, a Minkowski distance, a Mahalanobis distance, a Jaccard similarity coefficient, or any suitable similarity metric. As an example and not by way of limitation, a similarity metric of \vec{v}_1 and \vec{v}_2 may be a cosine similarity $\frac{\vec{v}_1 \cdot \vec{v}_2}{\|\vec{v}_1\| \|\vec{v}_2\|}$. As another example and not by way of limitation, a similarity metric of \vec{v}_1 and \vec{v}_2 may be a Euclidean distance $\|\vec{v}_1 - \vec{v}_2\|$. A similarity metric of two vectors may represent how similar the two objects or n-grams corresponding to the two vectors, respectively, are to one another, as measured by the distance between the two vectors in the vector space 800. As an example and not by way of limitation, vector 810 and vector 820 may correspond to objects that are more similar to one another than the objects corresponding to vector 810 and vector 880, based on the distance between the respective vectors. Although this disclosure describes calculating a similarity metric between vectors in a particular manner, this disclosure contemplates calculating a similarity metric between vectors in any suitable manner.

[113] More information on vector spaces, embeddings, feature vectors, and similarity metrics may be found in U.S. Patent Application No. 14/949436, filed 23 November 2015, U.S.

Patent Application No. 15/286315, filed 05 October 2016, and U.S. Patent Application No. 15/365789, filed 30 November 2016, each of which is incorporated by reference.

Artificial Neural Networks

[114] FIG. 9 illustrates an example artificial neural network (“ANN”) 900. In particular embodiments, an ANN may refer to a computational model comprising one or more nodes. Example ANN 900 may comprise an input layer 910, hidden layers 920, 930, 960, and an output layer 950. Each layer of the ANN 900 may comprise one or more nodes, such as a node 905 or a node 915. In particular embodiments, each node of an ANN may be connected to another node of the ANN. As an example and not by way of limitation, each node of the input layer 910 may be connected to one or more nodes of the hidden layer 920. In particular embodiments, one or more nodes may be a bias node (e.g., a node in a layer that is not connected to and does not receive input from any node in a previous layer). In particular embodiments, each node in each layer may be connected to one or more nodes of a previous or subsequent layer. Although FIG. 9 depicts a particular ANN with a particular number of layers, a particular number of nodes, and particular connections between nodes, this disclosure contemplates any suitable ANN with any suitable number of layers, any suitable number of nodes, and any suitable connections between nodes. As an example and not by way of limitation, although FIG. 9 depicts a connection between each node of the input layer 910 and each node of the hidden layer 920, one or more nodes of the input layer 910 may not be connected to one or more nodes of the hidden layer 920.

[115] In particular embodiments, an ANN may be a feedforward ANN (e.g., an ANN with no cycles or loops where communication between nodes flows in one direction beginning with the input layer and proceeding to successive layers). As an example and not by way of limitation, the input to each node of the hidden layer 920 may comprise the output of one or more nodes of the input layer 910. As another example and not by way of limitation, the input to each node of the output layer 950 may comprise the output of one or more nodes of the hidden layer 960. In particular embodiments, an ANN may be a deep neural network (e.g., a neural network comprising at least two hidden layers). In particular embodiments, an ANN may be a deep residual network. A deep residual network may be a feedforward ANN comprising hidden layers organized into residual blocks. The input into each residual block after the first residual block may be a function of the output of the previous residual block and the input of the previous residual block.

As an example and not by way of limitation, the input into residual block N may be $F(x) + x$, where $F(x)$ may be the output of residual block $N - 1$, x may be the input into residual block $N - 1$. Although this disclosure describes a particular ANN, this disclosure contemplates any suitable ANN.

[116] In particular embodiments, an activation function may correspond to each node of an ANN. An activation function of a node may define the output of a node for a given input. In particular embodiments, an input to a node may comprise a set of inputs. As an example and not by way of limitation, an activation function may be an identity function, a binary step function, a logistic function, or any other suitable function. As another example and not by way of limitation, an activation function for a node k may be the sigmoid function $F_k(s_k) = \frac{1}{1+e^{-s_k}}$, the hyperbolic tangent function $F_k(s_k) = \frac{e^{s_k} - e^{-s_k}}{e^{s_k} + e^{-s_k}}$, the rectifier $F_k(s_k) = \max(0, s_k)$, or any other suitable function $F_k(s_k)$, where s_k may be the effective input to node k . In particular embodiments, the input of an activation function corresponding to a node may be weighted. Each node may generate output using a corresponding activation function based on weighted inputs. In particular embodiments, each connection between nodes may be associated with a weight. As an example and not by way of limitation, a connection 925 between the node 905 and the node 915 may have a weighting coefficient of 0.4, which may indicate that 0.4 multiplied by the output of the node 905 is used as an input to the node 915. As another example and not by way of limitation, the output y_k of node k may be $y_k = F_k(s_k)$, where F_k may be the activation function corresponding to node k , $s_k = \sum_j (w_{jk} x_j)$ may be the effective input to node k , x_j may be the output of a node j connected to node k , and w_{jk} may be the weighting coefficient between node j and node k . In particular embodiments, the input to nodes of the input layer may be based on a vector representing an object. Although this disclosure describes particular inputs to and outputs of nodes, this disclosure contemplates any suitable inputs to and outputs of nodes. Moreover, although this disclosure may describe particular connections and weights between nodes, this disclosure contemplates any suitable connections and weights between nodes.

[117] In particular embodiments, an ANN may be trained using training data. As an example and not by way of limitation, training data may comprise inputs to the ANN 900 and an expected output. As another example and not by way of limitation, training data may comprise

vectors each representing a training object and an expected label for each training object. In particular embodiments, training an ANN may comprise modifying the weights associated with the connections between nodes of the ANN by optimizing an objective function. As an example and not by way of limitation, a training method may be used (e.g., the conjugate gradient method, the gradient descent method, the stochastic gradient descent) to backpropagate the sum-of-squares error measured as a distances between each vector representing a training object (e.g., using a cost function that minimizes the sum-of-squares error). In particular embodiments, an ANN may be trained using a dropout technique. As an example and not by way of limitation, one or more nodes may be temporarily omitted (e.g., receive no input and generate no output) while training. For each training object, one or more nodes of the ANN may have some probability of being omitted. The nodes that are omitted for a particular training object may be different than the nodes omitted for other training objects (e.g., the nodes may be temporarily omitted on an object-by-object basis). Although this disclosure describes training an ANN in a particular manner, this disclosure contemplates training an ANN in any suitable manner.

Privacy

[118] In particular embodiments, one or more objects (e.g., content or other types of objects) of a computing system may be associated with one or more privacy settings. The one or more objects may be stored on or otherwise associated with any suitable computing system or application, such as, for example, a social-networking system 160, a client system 130, an assistant system 140, a third-party system 170, a social-networking application, an assistant application, a messaging application, a photo-sharing application, or any other suitable computing system or application. Although the examples discussed herein are in the context of an online social network, these privacy settings may be applied to any other suitable computing system. Privacy settings (or “access settings”) for an object may be stored in any suitable manner, such as, for example, in association with the object, in an index on an authorization server, in another suitable manner, or any suitable combination thereof. A privacy setting for an object may specify how the object (or particular information associated with the object) can be accessed, stored, or otherwise used (e.g., viewed, shared, modified, copied, executed, surfaced, or identified) within the online social network. When privacy settings for an object allow a particular user or other entity to access that object, the object may be described as being “visible” with respect to that user or other entity. As

an example and not by way of limitation, a user of the online social network may specify privacy settings for a user-profile page that identify a set of users that may access work-experience information on the user-profile page, thus excluding other users from accessing that information.

[119] In particular embodiments, privacy settings for an object may specify a “blocked list” of users or other entities that should not be allowed to access certain information associated with the object. In particular embodiments, the blocked list may include third-party entities. The blocked list may specify one or more users or entities for which an object is not visible. As an example and not by way of limitation, a user may specify a set of users who may not access photo albums associated with the user, thus excluding those users from accessing the photo albums (while also possibly allowing certain users not within the specified set of users to access the photo albums). In particular embodiments, privacy settings may be associated with particular social-graph elements. Privacy settings of a social-graph element, such as a node or an edge, may specify how the social-graph element, information associated with the social-graph element, or objects associated with the social-graph element can be accessed using the online social network. As an example and not by way of limitation, a particular concept node 704 corresponding to a particular photo may have a privacy setting specifying that the photo may be accessed only by users tagged in the photo and friends of the users tagged in the photo. In particular embodiments, privacy settings may allow users to opt in to or opt out of having their content, information, or actions stored/logged by the social-networking system 160 or assistant system 140 or shared with other systems (e.g., a third-party system 170). Although this disclosure describes using particular privacy settings in a particular manner, this disclosure contemplates using any suitable privacy settings in any suitable manner.

[120] In particular embodiments, privacy settings may be based on one or more nodes or edges of a social graph 700. A privacy setting may be specified for one or more edges 706 or edge-types of the social graph 700, or with respect to one or more nodes 702, 704 or node-types of the social graph 700. The privacy settings applied to a particular edge 706 connecting two nodes may control whether the relationship between the two entities corresponding to the nodes is visible to other users of the online social network. Similarly, the privacy settings applied to a particular node may control whether the user or concept corresponding to the node is visible to other users of the online social network. As an example and not by way of limitation, a first user may share an object

to the social-networking system 160. The object may be associated with a concept node 704 connected to a user node 702 of the first user by an edge 706. The first user may specify privacy settings that apply to a particular edge 706 connecting to the concept node 704 of the object, or may specify privacy settings that apply to all edges 706 connecting to the concept node 704. As another example and not by way of limitation, the first user may share a set of objects of a particular object-type (e.g., a set of images). The first user may specify privacy settings with respect to all objects associated with the first user of that particular object-type as having a particular privacy setting (e.g., specifying that all images posted by the first user are visible only to friends of the first user and/or users tagged in the images).

[121] In particular embodiments, the social-networking system 160 may present a “privacy wizard” (e.g., within a webpage, a module, one or more dialog boxes, or any other suitable interface) to the first user to assist the first user in specifying one or more privacy settings. The privacy wizard may display instructions, suitable privacy-related information, current privacy settings, one or more input fields for accepting one or more inputs from the first user specifying a change or confirmation of privacy settings, or any suitable combination thereof. In particular embodiments, the social-networking system 160 may offer a “dashboard” functionality to the first user that may display, to the first user, current privacy settings of the first user. The dashboard functionality may be displayed to the first user at any appropriate time (e.g., following an input from the first user summoning the dashboard functionality, following the occurrence of a particular event or trigger action). The dashboard functionality may allow the first user to modify one or more of the first user’s current privacy settings at any time, in any suitable manner (e.g., redirecting the first user to the privacy wizard).

[122] Privacy settings associated with an object may specify any suitable granularity of permitted access or denial of access. As an example and not by way of limitation, access or denial of access may be specified for particular users (e.g., only me, my roommates, my boss), users within a particular degree-of-separation (e.g., friends, friends-of-friends), user groups (e.g., the gaming club, my family), user networks (e.g., employees of particular employers, students or alumni of particular university), all users (“public”), no users (“private”), users of third-party systems 170, particular applications (e.g., third-party applications, external websites), other suitable entities, or any suitable combination thereof. Although this disclosure describes particular

granularities of permitted access or denial of access, this disclosure contemplates any suitable granularities of permitted access or denial of access.

[123] In particular embodiments, one or more servers 162 may be authorization/privacy servers for enforcing privacy settings. In response to a request from a user (or other entity) for a particular object stored in a data store 164, the social-networking system 160 may send a request to the data store 164 for the object. The request may identify the user associated with the request and the object may be sent only to the user (or a client system 130 of the user) if the authorization server determines that the user is authorized to access the object based on the privacy settings associated with the object. If the requesting user is not authorized to access the object, the authorization server may prevent the requested object from being retrieved from the data store 164 or may prevent the requested object from being sent to the user. In the search-query context, an object may be provided as a search result only if the querying user is authorized to access the object, e.g., if the privacy settings for the object allow it to be surfaced to, discovered by, or otherwise visible to the querying user. In particular embodiments, an object may represent content that is visible to a user through a newsfeed of the user. As an example and not by way of limitation, one or more objects may be visible to a user's "Trending" page. In particular embodiments, an object may correspond to a particular user. The object may be content associated with the particular user, or may be the particular user's account or information stored on the social-networking system 160, or other computing system. As an example and not by way of limitation, a first user may view one or more second users of an online social network through a "People You May Know" function of the online social network, or by viewing a list of friends of the first user. As an example and not by way of limitation, a first user may specify that they do not wish to see objects associated with a particular second user in their newsfeed or friends list. If the privacy settings for the object do not allow it to be surfaced to, discovered by, or visible to the user, the object may be excluded from the search results. Although this disclosure describes enforcing privacy settings in a particular manner, this disclosure contemplates enforcing privacy settings in any suitable manner.

[124] In particular embodiments, different objects of the same type associated with a user may have different privacy settings. Different types of objects associated with a user may have different types of privacy settings. As an example and not by way of limitation, a first user may specify that the first user's status updates are public, but any images shared by the first user are

visible only to the first user's friends on the online social network. As another example and not by way of limitation, a user may specify different privacy settings for different types of entities, such as individual users, friends-of-friends, followers, user groups, or corporate entities. As another example and not by way of limitation, a first user may specify a group of users that may view videos posted by the first user, while keeping the videos from being visible to the first user's employer. In particular embodiments, different privacy settings may be provided for different user groups or user demographics. As an example and not by way of limitation, a first user may specify that other users who attend the same university as the first user may view the first user's pictures, but that other users who are family members of the first user may not view those same pictures.

[125] In particular embodiments, the social-networking system 160 may provide one or more default privacy settings for each object of a particular object-type. A privacy setting for an object that is set to a default may be changed by a user associated with that object. As an example and not by way of limitation, all images posted by a first user may have a default privacy setting of being visible only to friends of the first user and, for a particular image, the first user may change the privacy setting for the image to be visible to friends and friends-of-friends.

[126] In particular embodiments, privacy settings may allow a first user to specify (e.g., by opting out, by not opting in) whether the social-networking system 160 or assistant system 140 may receive, collect, log, or store particular objects or information associated with the user for any purpose. In particular embodiments, privacy settings may allow the first user to specify whether particular applications or processes may access, store, or use particular objects or information associated with the user. The privacy settings may allow the first user to opt in or opt out of having objects or information accessed, stored, or used by specific applications or processes. The social-networking system 160 or assistant system 140 may access such information in order to provide a particular function or service to the first user, without the social-networking system 160 or assistant system 140 having access to that information for any other purposes. Before accessing, storing, or using such objects or information, the social-networking system 160 or assistant system 140 may prompt the user to provide privacy settings specifying which applications or processes, if any, may access, store, or use the object or information prior to allowing any such action. As an example and not by way of limitation, a first user may transmit a message to a second user via an application

related to the online social network (e.g., a messaging app), and may specify privacy settings that such messages should not be stored by the social-networking system 160 or assistant system 140.

[127] In particular embodiments, a user may specify whether particular types of objects or information associated with the first user may be accessed, stored, or used by the social-networking system 160 or assistant system 140. As an example and not by way of limitation, the first user may specify that images sent by the first user through the social-networking system 160 or assistant system 140 may not be stored by the social-networking system 160 or assistant system 140. As another example and not by way of limitation, a first user may specify that messages sent from the first user to a particular second user may not be stored by the social-networking system 160 or assistant system 140. As yet another example and not by way of limitation, a first user may specify that all objects sent via a particular application may be saved by the social-networking system 160 or assistant system 140.

[128] In particular embodiments, privacy settings may allow a first user to specify whether particular objects or information associated with the first user may be accessed from particular client systems 130 or third-party systems 170. The privacy settings may allow the first user to opt in or opt out of having objects or information accessed from a particular device (e.g., the phone book on a user's smart phone), from a particular application (e.g., a messaging app), or from a particular system (e.g., an email server). The social-networking system 160 or assistant system 140 may provide default privacy settings with respect to each device, system, or application, and/or the first user may be prompted to specify a particular privacy setting for each context. As an example and not by way of limitation, the first user may utilize a location-services feature of the social-networking system 160 or assistant system 140 to provide recommendations for restaurants or other places in proximity to the user. The first user's default privacy settings may specify that the social-networking system 160 or assistant system 140 may use location information provided from a client device 130 of the first user to provide the location-based services, but that the social-networking system 160 or assistant system 140 may not store the location information of the first user or provide it to any third-party system 170. The first user may then update the privacy settings to allow location information to be used by a third-party image-sharing application in order to geo-tag photos.

Privacy Settings Based on Location

[129] In particular embodiments, privacy settings may allow a user to specify one or more geographic locations from which objects can be accessed. Access or denial of access to the objects may depend on the geographic location of a user who is attempting to access the objects. As an example and not by way of limitation, a user may share an object and specify that only users in the same city may access or view the object. As another example and not by way of limitation, a first user may share an object and specify that the object is visible to second users only while the first user is in a particular location. If the first user leaves the particular location, the object may no longer be visible to the second users. As another example and not by way of limitation, a first user may specify that an object is visible only to second users within a threshold distance from the first user. If the first user subsequently changes location, the original second users with access to the object may lose access, while a new group of second users may gain access as they come within the threshold distance of the first user.

Privacy Settings for User-Authentication and Experience-Personalization Information

[130] In particular embodiments, the social-networking system 160 or assistant system 140 may have functionalities that may use, as inputs, personal or biometric information of a user for user-authentication or experience-personalization purposes. A user may opt to make use of these functionalities to enhance their experience on the online social network. As an example and not by way of limitation, a user may provide personal or biometric information to the social-networking system 160 or assistant system 140. The user's privacy settings may specify that such information may be used only for particular processes, such as authentication, and further specify that such information may not be shared with any third-party system 170 or used for other processes or applications associated with the social-networking system 160 or assistant system 140. As another example and not by way of limitation, the social-networking system 160 may provide a functionality for a user to provide voice-print recordings to the online social network. As an example and not by way of limitation, if a user wishes to utilize this function of the online social network, the user may provide a voice recording of his or her own voice to provide a status update on the online social network. The recording of the voice-input may be compared to a voice print of the user to determine what words were spoken by the user. The user's privacy setting may

specify that such voice recording may be used only for voice-input purposes (e.g., to authenticate the user, to send voice messages, to improve voice recognition in order to use voice-operated features of the online social network), and further specify that such voice recording may not be shared with any third-party system 170 or used by other processes or applications associated with the social-networking system 160. As another example and not by way of limitation, the social-networking system 160 may provide a functionality for a user to provide a reference image (e.g., a facial profile, a retinal scan) to the online social network. The online social network may compare the reference image against a later-received image input (e.g., to authenticate the user, to tag the user in photos). The user's privacy setting may specify that such voice recording may be used only for a limited purpose (e.g., authentication, tagging the user in photos), and further specify that such voice recording may not be shared with any third-party system 170 or used by other processes or applications associated with the social-networking system 160.

Systems and Methods

[131] FIG. 10 illustrates an example computer system 1000. In particular embodiments, one or more computer systems 1000 perform one or more steps of one or more methods described or illustrated herein. In particular embodiments, one or more computer systems 1000 provide functionality described or illustrated herein. In particular embodiments, software running on one or more computer systems 1000 performs one or more steps of one or more methods described or illustrated herein or provides functionality described or illustrated herein. Particular embodiments include one or more portions of one or more computer systems 1000. Herein, reference to a computer system may encompass a computing device, and vice versa, where appropriate. Moreover, reference to a computer system may encompass one or more computer systems, where appropriate.

[132] This disclosure contemplates any suitable number of computer systems 1000. This disclosure contemplates computer system 1000 taking any suitable physical form. As example and not by way of limitation, computer system 1000 may be an embedded computer system, a system-on-chip (SOC), a single-board computer system (SBC) (such as, for example, a computer-on-module (COM) or system-on-module (SOM)), a desktop computer system, a laptop or notebook computer system, an interactive kiosk, a mainframe, a mesh of computer systems, a mobile telephone, a personal digital assistant (PDA), a server, a tablet computer system, or a combination

of two or more of these. Where appropriate, computer system 1000 may include one or more computer systems 1000; be unitary or distributed; span multiple locations; span multiple machines; span multiple data centers; or reside in a cloud, which may include one or more cloud components in one or more networks. Where appropriate, one or more computer systems 1000 may perform without substantial spatial or temporal limitation one or more steps of one or more methods described or illustrated herein. As an example and not by way of limitation, one or more computer systems 1000 may perform in real time or in batch mode one or more steps of one or more methods described or illustrated herein. One or more computer systems 1000 may perform at different times or at different locations one or more steps of one or more methods described or illustrated herein, where appropriate.

[133] In particular embodiments, computer system 1000 includes a processor 1002, memory 1004, storage 1006, an input/output (I/O) interface 1008, a communication interface 1010, and a bus 1012. Although this disclosure describes and illustrates a particular computer system having a particular number of particular components in a particular arrangement, this disclosure contemplates any suitable computer system having any suitable number of any suitable components in any suitable arrangement.

[134] In particular embodiments, processor 1002 includes hardware for executing instructions, such as those making up a computer program. As an example and not by way of limitation, to execute instructions, processor 1002 may retrieve (or fetch) the instructions from an internal register, an internal cache, memory 1004, or storage 1006; decode and execute them; and then write one or more results to an internal register, an internal cache, memory 1004, or storage 1006. In particular embodiments, processor 1002 may include one or more internal caches for data, instructions, or addresses. This disclosure contemplates processor 1002 including any suitable number of any suitable internal caches, where appropriate. As an example and not by way of limitation, processor 1002 may include one or more instruction caches, one or more data caches, and one or more translation lookaside buffers (TLBs). Instructions in the instruction caches may be copies of instructions in memory 1004 or storage 1006, and the instruction caches may speed up retrieval of those instructions by processor 1002. Data in the data caches may be copies of data in memory 1004 or storage 1006 for instructions executing at processor 1002 to operate on; the results of previous instructions executed at processor 1002 for access by subsequent instructions

executing at processor 1002 or for writing to memory 1004 or storage 1006; or other suitable data. The data caches may speed up read or write operations by processor 1002. The TLBs may speed up virtual-address translation for processor 1002. In particular embodiments, processor 1002 may include one or more internal registers for data, instructions, or addresses. This disclosure contemplates processor 1002 including any suitable number of any suitable internal registers, where appropriate. Where appropriate, processor 1002 may include one or more arithmetic logic units (ALUs); be a multi-core processor; or include one or more processors 1002. Although this disclosure describes and illustrates a particular processor, this disclosure contemplates any suitable processor.

[135] In particular embodiments, memory 1004 includes main memory for storing instructions for processor 1002 to execute or data for processor 1002 to operate on. As an example and not by way of limitation, computer system 1000 may load instructions from storage 1006 or another source (such as, for example, another computer system 1000) to memory 1004. Processor 1002 may then load the instructions from memory 1004 to an internal register or internal cache. To execute the instructions, processor 1002 may retrieve the instructions from the internal register or internal cache and decode them. During or after execution of the instructions, processor 1002 may write one or more results (which may be intermediate or final results) to the internal register or internal cache. Processor 1002 may then write one or more of those results to memory 1004. In particular embodiments, processor 1002 executes only instructions in one or more internal registers or internal caches or in memory 1004 (as opposed to storage 1006 or elsewhere) and operates only on data in one or more internal registers or internal caches or in memory 1004 (as opposed to storage 1006 or elsewhere). One or more memory buses (which may each include an address bus and a data bus) may couple processor 1002 to memory 1004. Bus 1012 may include one or more memory buses, as described below. In particular embodiments, one or more memory management units (MMUs) reside between processor 1002 and memory 1004 and facilitate accesses to memory 1004 requested by processor 1002. In particular embodiments, memory 1004 includes random access memory (RAM). This RAM may be volatile memory, where appropriate. Where appropriate, this RAM may be dynamic RAM (DRAM) or static RAM (SRAM). Moreover, where appropriate, this RAM may be single-ported or multi-ported RAM. This disclosure contemplates any suitable RAM. Memory 1004 may include one or more memories 1004, where appropriate.

Although this disclosure describes and illustrates particular memory, this disclosure contemplates any suitable memory.

[136] In particular embodiments, storage 1006 includes mass storage for data or instructions. As an example and not by way of limitation, storage 1006 may include a hard disk drive (HDD), a floppy disk drive, flash memory, an optical disc, a magneto-optical disc, magnetic tape, or a Universal Serial Bus (USB) drive or a combination of two or more of these. Storage 1006 may include removable or non-removable (or fixed) media, where appropriate. Storage 1006 may be internal or external to computer system 1000, where appropriate. In particular embodiments, storage 1006 is non-volatile, solid-state memory. In particular embodiments, storage 1006 includes read-only memory (ROM). Where appropriate, this ROM may be mask-programmed ROM, programmable ROM (PROM), erasable PROM (EPROM), electrically erasable PROM (EEPROM), electrically alterable ROM (EAROM), or flash memory or a combination of two or more of these. This disclosure contemplates mass storage 1006 taking any suitable physical form. Storage 1006 may include one or more storage control units facilitating communication between processor 1002 and storage 1006, where appropriate. Where appropriate, storage 1006 may include one or more storages 1006. Although this disclosure describes and illustrates particular storage, this disclosure contemplates any suitable storage.

[137] In particular embodiments, I/O interface 1008 includes hardware, software, or both, providing one or more interfaces for communication between computer system 1000 and one or more I/O devices. Computer system 1000 may include one or more of these I/O devices, where appropriate. One or more of these I/O devices may enable communication between a person and computer system 1000. As an example and not by way of limitation, an I/O device may include a keyboard, keypad, microphone, monitor, mouse, printer, scanner, speaker, still camera, stylus, tablet, touch screen, trackball, video camera, another suitable I/O device or a combination of two or more of these. An I/O device may include one or more sensors. This disclosure contemplates any suitable I/O devices and any suitable I/O interfaces 1008 for them. Where appropriate, I/O interface 1008 may include one or more device or software drivers enabling processor 1002 to drive one or more of these I/O devices. I/O interface 1008 may include one or more I/O interfaces 1008, where appropriate. Although this disclosure describes and illustrates a particular I/O interface, this disclosure contemplates any suitable I/O interface.

[138] In particular embodiments, communication interface 1010 includes hardware, software, or both providing one or more interfaces for communication (such as, for example, packet-based communication) between computer system 1000 and one or more other computer systems 1000 or one or more networks. As an example and not by way of limitation, communication interface 1010 may include a network interface controller (NIC) or network adapter for communicating with an Ethernet or other wire-based network or a wireless NIC (WNIC) or wireless adapter for communicating with a wireless network, such as a WI-FI network. This disclosure contemplates any suitable network and any suitable communication interface 1010 for it. As an example and not by way of limitation, computer system 1000 may communicate with an ad hoc network, a personal area network (PAN), a local area network (LAN), a wide area network (WAN), a metropolitan area network (MAN), or one or more portions of the Internet or a combination of two or more of these. One or more portions of one or more of these networks may be wired or wireless. As an example, computer system 1000 may communicate with a wireless PAN (WPAN) (such as, for example, a BLUETOOTH WPAN), a WI-FI network, a WI-MAX network, a cellular telephone network (such as, for example, a Global System for Mobile Communications (GSM) network), or other suitable wireless network or a combination of two or more of these. Computer system 1000 may include any suitable communication interface 1010 for any of these networks, where appropriate. Communication interface 1010 may include one or more communication interfaces 1010, where appropriate. Although this disclosure describes and illustrates a particular communication interface, this disclosure contemplates any suitable communication interface.

[139] In particular embodiments, bus 1012 includes hardware, software, or both coupling components of computer system 1000 to each other. As an example and not by way of limitation, bus 1012 may include an Accelerated Graphics Port (AGP) or other graphics bus, an Enhanced Industry Standard Architecture (EISA) bus, a front-side bus (FSB), a HYPERTRANSPORT (HT) interconnect, an Industry Standard Architecture (ISA) bus, an INFINIBAND interconnect, a low-pin-count (LPC) bus, a memory bus, a Micro Channel Architecture (MCA) bus, a Peripheral Component Interconnect (PCI) bus, a PCI-Express (PCIe) bus, a serial advanced technology attachment (SATA) bus, a Video Electronics Standards Association local (VLB) bus, or another suitable bus or a combination of two or more of these. Bus 1012 may include one or more buses

1012, where appropriate. Although this disclosure describes and illustrates a particular bus, this disclosure contemplates any suitable bus or interconnect.

[140] Herein, a computer-readable non-transitory storage medium or media may include one or more semiconductor-based or other integrated circuits (ICs) (such, as for example, field-programmable gate arrays (FPGAs) or application-specific ICs (ASICs)), hard disk drives (HDDs), hybrid hard drives (HHDs), optical discs, optical disc drives (ODDs), magneto-optical discs, magneto-optical drives, floppy diskettes, floppy disk drives (FDDs), magnetic tapes, solid-state drives (SSDs), RAM-drives, SECURE DIGITAL cards or drives, any other suitable computer-readable non-transitory storage media, or any suitable combination of two or more of these, where appropriate. A computer-readable non-transitory storage medium may be volatile, non-volatile, or a combination of volatile and non-volatile, where appropriate.

Miscellaneous

[141] Herein, “or” is inclusive and not exclusive, unless expressly indicated otherwise or indicated otherwise by context. Therefore, herein, “A or B” means “A, B, or both,” unless expressly indicated otherwise or indicated otherwise by context. Moreover, “and” is both joint and several, unless expressly indicated otherwise or indicated otherwise by context. Therefore, herein, “A and B” means “A and B, jointly or severally,” unless expressly indicated otherwise or indicated otherwise by context.

[142] The scope of this disclosure encompasses all changes, substitutions, variations, alterations, and modifications to the example embodiments described or illustrated herein that a person having ordinary skill in the art would comprehend. The scope of this disclosure is not limited to the example embodiments described or illustrated herein. Moreover, although this disclosure describes and illustrates respective embodiments herein as including particular components, elements, feature, functions, operations, or steps, any of these embodiments may include any combination or permutation of any of the components, elements, features, functions, operations, or steps described or illustrated anywhere herein that a person having ordinary skill in the art would comprehend. Furthermore, reference in the appended claims to an apparatus or system or a component of an apparatus or system being adapted to, arranged to, capable of, configured to, enabled to, operable to, or operative to perform a particular function encompasses that apparatus, system, component, whether or not it or that particular function is activated, turned

on, or unlocked, as long as that apparatus, system, or component is so adapted, arranged, capable, configured, enabled, operable, or operative. Additionally, although this disclosure describes or illustrates particular embodiments as providing particular advantages, particular embodiments may provide none, some, or all of these advantages.

CLAIMS

What is claimed is:

1. A method comprising, by one or more computing systems:
 - accessing, from an online social network, a plurality of content objects associated with a first user;
 - accessing a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;
 - accessing conversational data associated with the first user;
 - determining one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and
 - generating a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

2. The method of Claim 1, wherein the plurality of content objects comprise one or more of:
 - news feed posts;
 - news feed comments; or
 - search history data.

3. The method of Claim 1, wherein the accessed content objects were created within a pre-specified time window.

4. The method of Claim 1, wherein the conversational data comprises one or more messages created within a pre-specified time window.

5. The method of Claim 1, wherein determining the one or more subjects associated with the first user is based on one or more machine-learning models, wherein the one or more machine-

learning models calculate one or more probabilities of the conversational data being associated with the one or more subjects, respectively.

6. The method of Claim 1, further comprising determining one or more sentiment indications associated with the conversational data, wherein the determining is based on one or more long-short term memory (LSTM) networks, and wherein the one or more LSTM networks calculate one or more probabilities of the conversational data being associated with the one or more sentiment indications, respectively.

7. The method of Claim 1, further comprising determining one or more domains associated with the conversational data, wherein the determining is based on one or more machine-learning models, and wherein the one or more machine-learning models calculate one or more probabilities of the conversational data being associated with the one or more domains, respectively.

8. The method of Claim 1, wherein the conversation data comprises a plurality of messages, and wherein the method further comprises processing the conversational data based on a sliding window analysis over the plurality of messages.

9. The method of Claim 1, further comprising determining contextual information associated with the first user based on the conversational data.

10. The method of Claim 9, wherein determining the contextual information associated with the first user comprises one or more of:

processing meta-data associated with the conversational data based on a machine-learning model, wherein the machine-learning model calculates a probability of the meta-data being associated with the contextual information; or

processing the conversational data based on the machine-learning model, wherein the machine-learning model calculates a probability of the conversational data being associated with the contextual information.

11. The method of Claim 9, wherein the contextual information comprises one or more of location information, time stamps, or social connections.

12. The method of Claim 1, wherein the customized user profile is based on a tree structure.

13. The method of Claim 12, wherein the tree structure comprises one or more layers.

14. The method of Claim 13, wherein the one or more layers are organized in a hierarchical structure comprising:

- (1) a topic layer comprising one or more topics;
- (2) a category layer comprising one or more categories associated with each topic of the topic layer; and
- (3) an entity layer comprising one or more entities associated with each category of the category layer.

15. The method of Claim 1, further comprising calculating the one or more confidence scores for the one or more subjects based on a frequency of the respective subject in the plurality of content objects and the conversational data.

16. The method of Claim 1, wherein each of the one or more subjects comprises a topic, a category, or an entity.

17. The method of Claim 1, further comprising determining entity linking information based on the conversational data.

18. One or more computer-readable non-transitory storage media embodying software that is operable when executed to:

access, from an online social network, a plurality of content objects associated with a first user;

access a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;

access conversational data associated with the first user;

determine one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and

generate a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

19. A system comprising: one or more processors; and a non-transitory memory coupled to the processors comprising instructions executable by the processors, the processors operable when executing the instructions to:

access, from an online social network, a plurality of content objects associated with a first user;

access a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;

access conversational data associated with the first user;

determine one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and

generate a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

20. A method, in particular for use in an assistant system for assisting a user to obtain information or services by enabling the user to interact with the assistant system with user input in conversations to get assistance, wherein the user input includes voice, text, image or video or any

combination of them, the assistant system in particular being enabled by the combination of computing devices, application programming interfaces (APIs), and the proliferation of applications on user devices,

the method comprising, by one or more computing systems:

accessing, from an online social network, a plurality of content objects associated with a first user;

accessing a baseline profile, wherein the baseline profile is based on ontology data from one or more information graphs;

accessing conversational data associated with the first user;

determining one or more subjects associated with the first user based on the plurality of content objects associated with the first user and the conversational data associated with the first user; and

generating a customized user profile for the first user based on the baseline profile, wherein the user profile comprises one or more confidence scores associated with the respective one or more subjects associated with the first user, wherein the one or more confidence scores are calculated based on the plurality of content objects associated with the first user and the conversational data associated with the first user.

21. The method of Claim 20, wherein the plurality of content objects comprise one or more of:
news feed posts;
news feed comments; or
search history data.

22. The method of Claim 20 or 21, wherein the accessed content objects were created within a pre-specified time window.

23. The method of any of Claims 20 to 22, wherein the conversational data comprises one or more messages created within a pre-specified time window.

24. The method of any of Claims 20 to 23, wherein determining the one or more subjects associated with the first user is based on one or more machine-learning models, wherein the one

or more machine-learning models calculate one or more probabilities of the conversational data being associated with the one or more subjects, respectively, wherein in particular support vector machines (SVM), a regression model, or a deep convolutional neural network (DCNN) is used.

25. The method of any of Claims 20 to 24, further comprising determining one or more sentiment indications associated with the conversational data, wherein the determining is based on one or more long-short term memory (LSTM) networks, and wherein the one or more LSTM networks calculate one or more probabilities of the conversational data being associated with the one or more sentiment indications, respectively.

26. The method of any of Claims 20 to 25, further comprising determining one or more domains associated with the conversational data, wherein the determining is based on one or more machine-learning models, and wherein the one or more machine-learning models calculate one or more probabilities of the conversational data being associated with the one or more domains, respectively.

27. The method of any of Claims 20 to 26, wherein the conversation data comprises a plurality of messages, and wherein the method further comprises processing the conversational data based on a sliding window analysis over the plurality of messages.

28. The method of any of Claims 20 to 27, further comprising determining contextual information associated with the first user based on the conversational data.

29. The method of Claim 28, wherein determining the contextual information associated with the first user comprises one or more of:

processing meta-data associated with the conversational data based on a machine-learning model, wherein the machine-learning model calculates a probability of the meta-data being associated with the contextual information; or

processing the conversational data based on the machine-learning model, wherein the machine-learning model calculates a probability of the conversational data being associated with the contextual information.

30. The method of Claim 28 or 29, wherein the contextual information comprises one or more of location information, time stamps, or social connections.

31. The method of any of Claims 20 to 30, wherein the customized user profile is based on a tree structure.

32. The method of Claim 31, wherein the tree structure comprises one or more layers.

33. The method of Claim 32, wherein the one or more layers are organized in a hierarchical structure comprising:

- (1) a topic layer comprising one or more topics;
- (2) a category layer comprising one or more categories associated with each topic of the topic layer; and
- (3) an entity layer comprising one or more entities associated with each category of the category layer.

34. The method of any of Claims 20 to 33, further comprising calculating the one or more confidence scores for the one or more subjects based on a frequency of the respective subject in the plurality of content objects and the conversational data.

35. The method of any of Claims 20 to 34, wherein each of the one or more subjects comprises a topic, a category, or an entity.

36. The method any of Claims 20 to 35, further comprising determining entity linking information based on the conversational data.

37. One or more computer-readable non-transitory storage media embodying software that is operable when executed to perform a method according to any of claims 20 to 36:

38. An assistant system for assisting a user to obtain information or services by enabling the user to interact with the assistant system with user input in conversations to get assistance, wherein

the user input includes voice, text, image or video or any combination of them, the assistant system in particular being enabled by the combination of computing devices, application programming interfaces (APIs), and the proliferation of applications on user devices,

the assistant system comprising:

one or more processors; and

a non-transitory memory coupled to the processors comprising instructions executable by the processors,

the processors operable when executing the instructions to perform a method according to any of claims 20 to 36.

39. The assistant system of claim 38 for assisting the user by executing at least one or more of the following features or steps:

create and store a user profile comprising both personal and contextual information associated with the user

analyze the user input using natural-language understanding, wherein the analysis may be based on the user profile for more personalized and context-aware understanding

resolve entities associated with the user input based on the analysis

interact with different agents to obtain information or services that are associated with the resolved entities

generate a response for the user regarding the information or services by using natural-language generation

through the interaction with the user, use dialog management techniques to manage and forward the conversation flow with the user

assist the user to effectively and efficiently digest the obtained information by summarizing the information

assist the user to be more engaging with an online social network by providing tools that help the user interact with the online social network (e.g., creating posts, comments, messages)

assist the user to manage different tasks such as keeping track of events

proactively execute pre-authorized tasks that are relevant to user interests and preferences based on the user profile, at a time relevant for the user, without a user input

check privacy settings whenever it is necessary to guarantee that accessing user profile and executing different tasks are subject to the user's privacy settings

automatically identify a user's intent and language register based on a user input, generate a communication content suitable for such intent and language register for the user, detect a change of the user's intent and language register, and dynamically adjust the generation of the communication content to fit the changed intent and language register.

40. The assistant system of Claim 38 or Claim 39, comprising at least one or more of the following components:

a messaging platform to receive a user input based on a text modality from the client system associated with the user and/or to receive user input based on an image or video modality and to process it using optical character recognition techniques within the messaging platform to convert the user input into text,

an audio speech recognition (ASR) module to receive a user input based on an audio modality (e.g., the user may speak to or send a video including speech) from the client system associated with the user and to convert the user input based on the audio modality into text,

an assistant xbot to receive the output of the messaging platform or the ASR module.

41. The assistant system of any of Claims 38 to 40, further comprising a user context engine (USE) (225), wherein the user context engine (225) stores the user profile of the user and wherein the user context engine (225) accesses the plurality of content objects (401) associated with the first user from the social-networking system or online social network (160).

42. The assistant system according to any of Claims 38 to 40, further comprising a dialog engine (235),

wherein the dialog engine (235) manages the dialog state and flow of the conversation between the user and the assistant xbot (215),

43. A system comprising

at least one client system (130), in particular an electronic device

at least one assistant system (140) according to any of claims 38 to 42,

connected to each other, in particular by a network (110),
wherein the client system comprises an assistant application (136) for allowing a user at the client system (130) to interact with the assistant system (140),
wherein the assistant application (136) communicates user input to the assistant system (140) and, based on the user input, the assistant system (140) generates responses and sends the generated responses to the assistant application (136) and the assistant application (136) presents the responses to the user at the client system (130),
wherein in particular the user input is audio or verbal and the response may be in text or also audio or verbal.

44. The system of claim 43 further comprising a social-networking system (160),
wherein the client system comprises in particular a social-networking application (134) for accessing the social networking system (160).

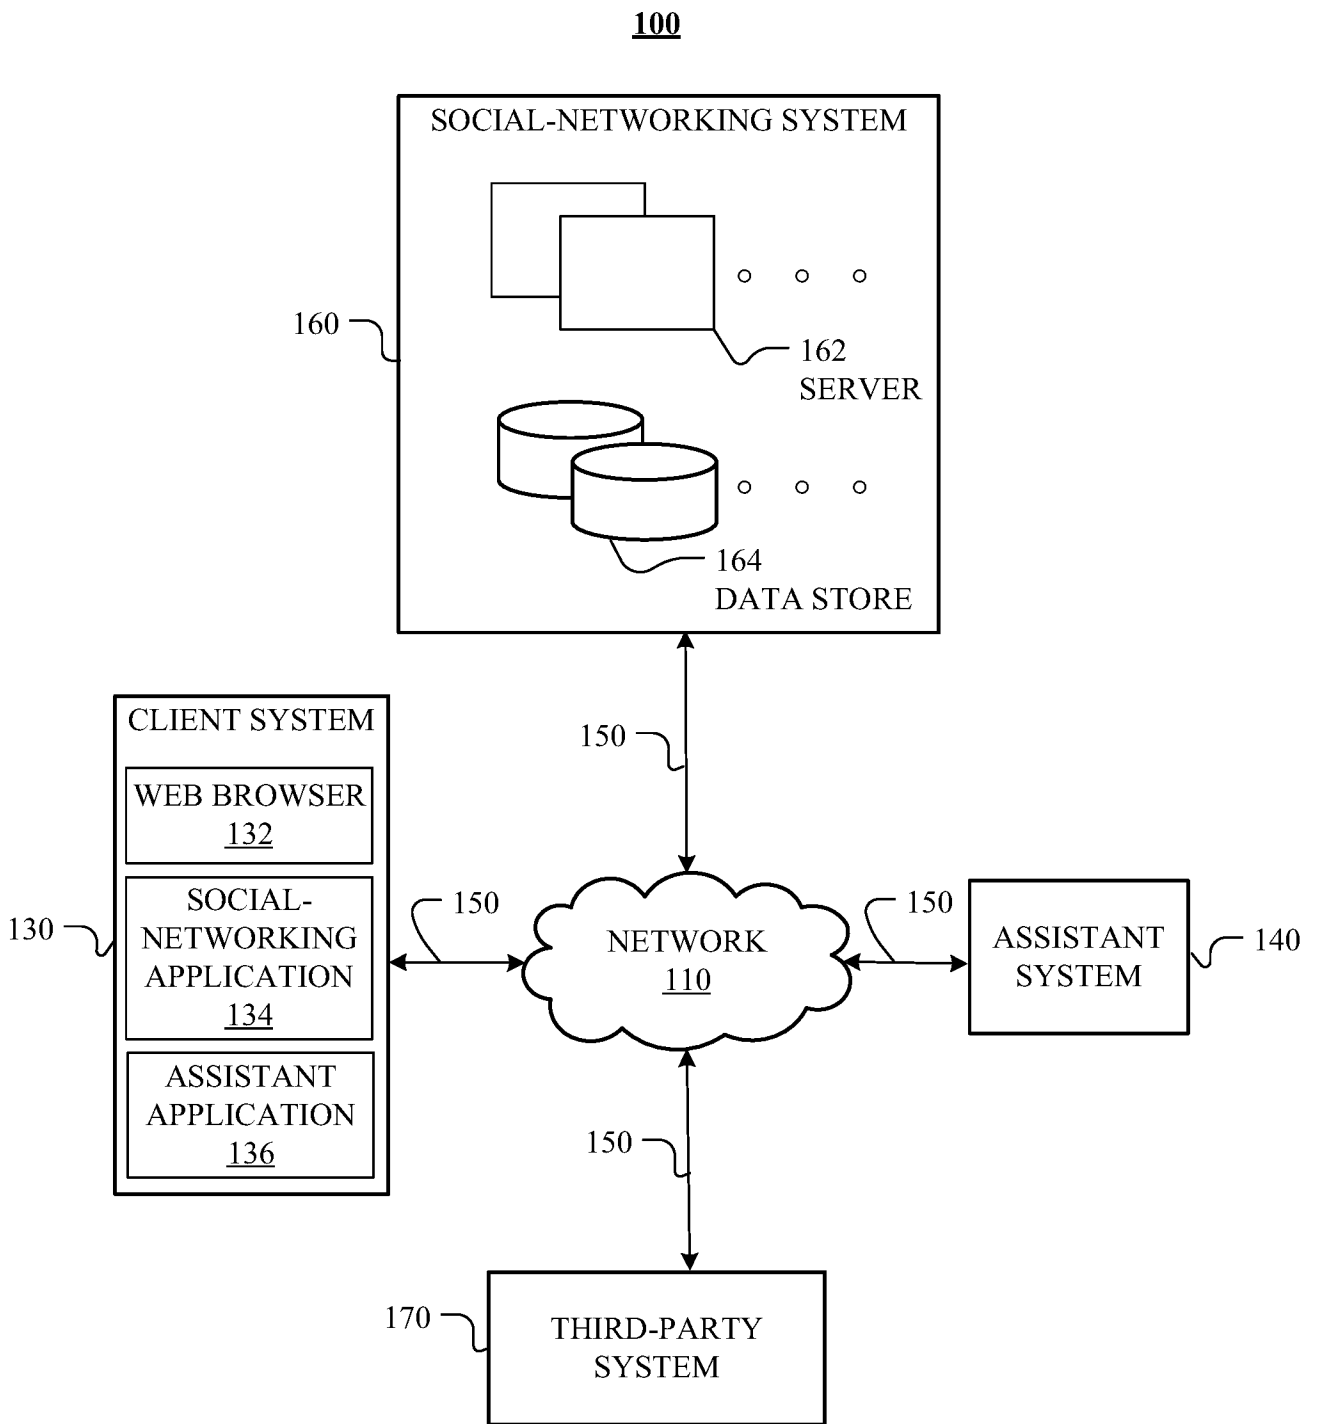


FIG. 1

200

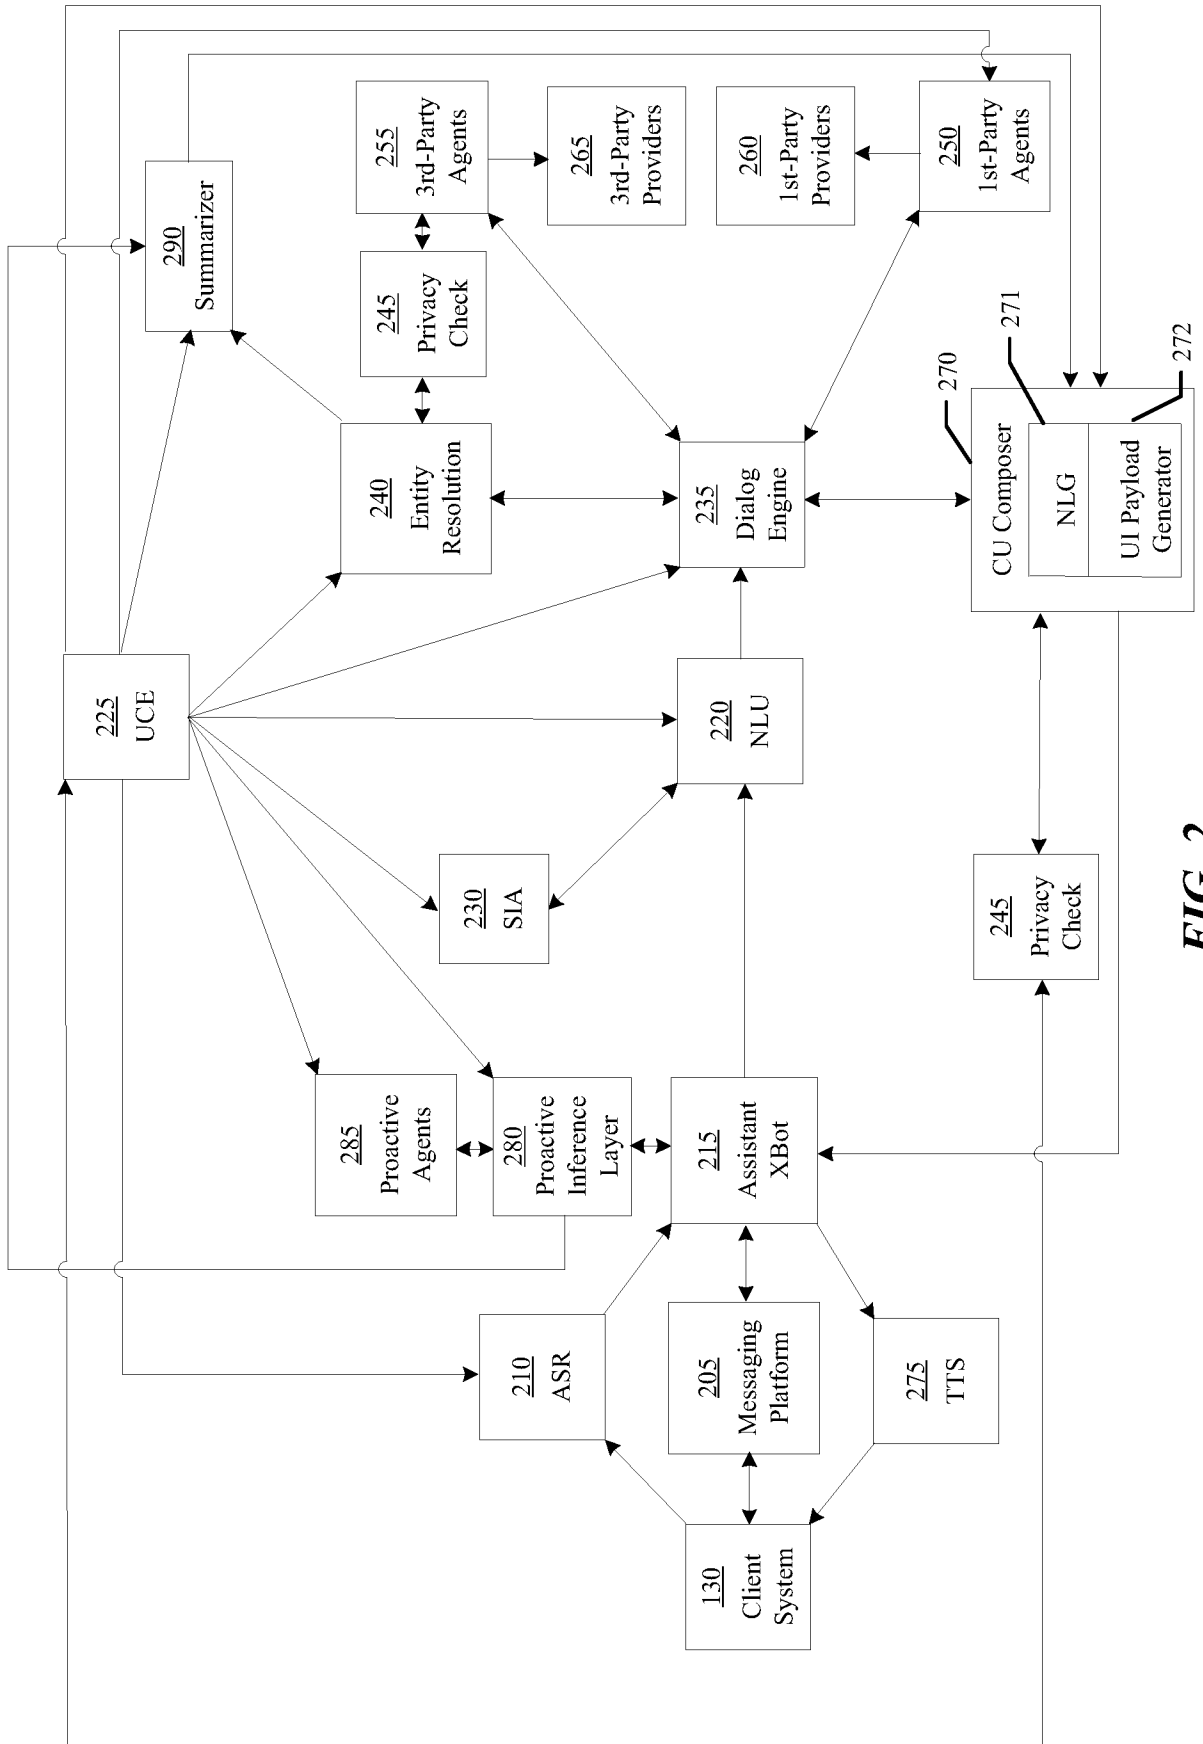


FIG. 2

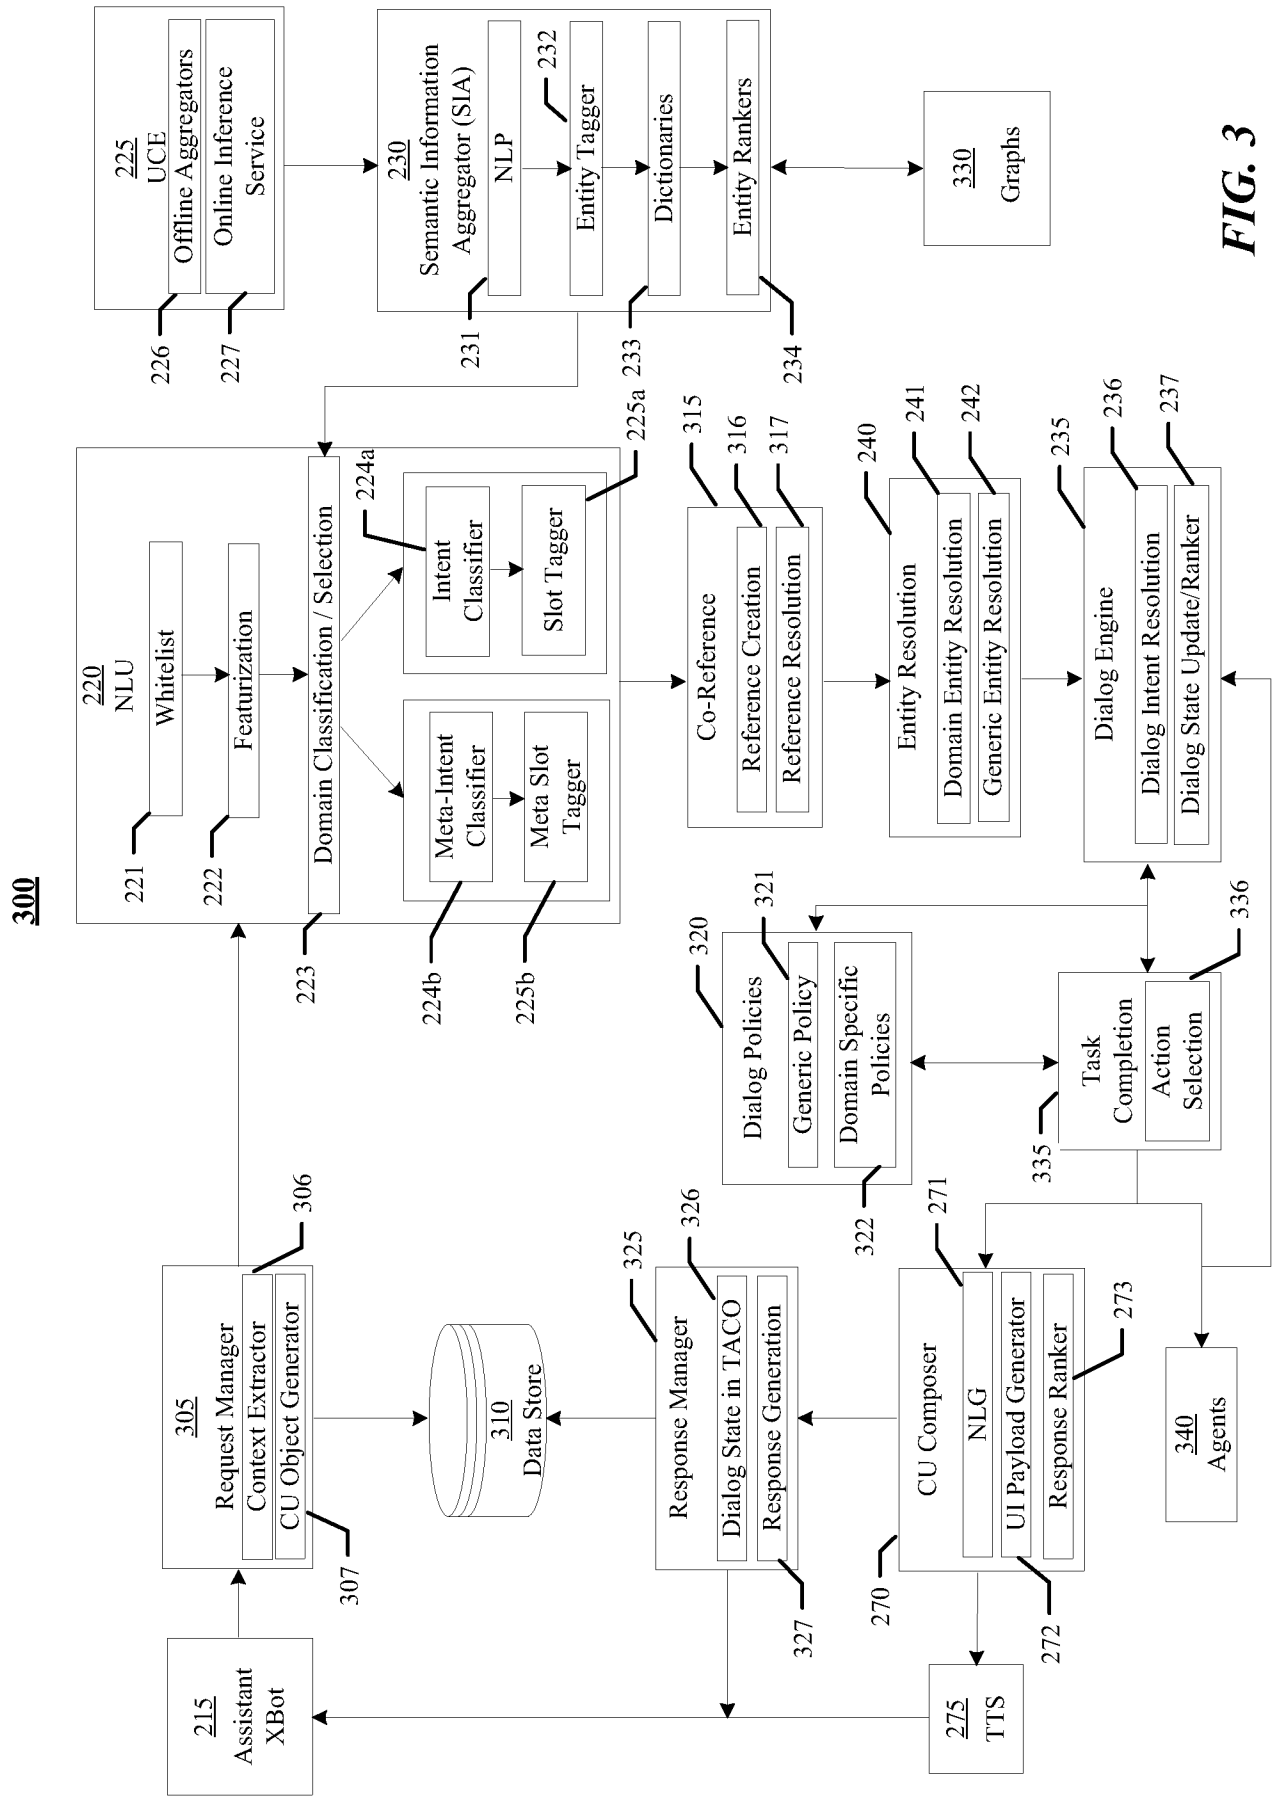


FIG. 3

400

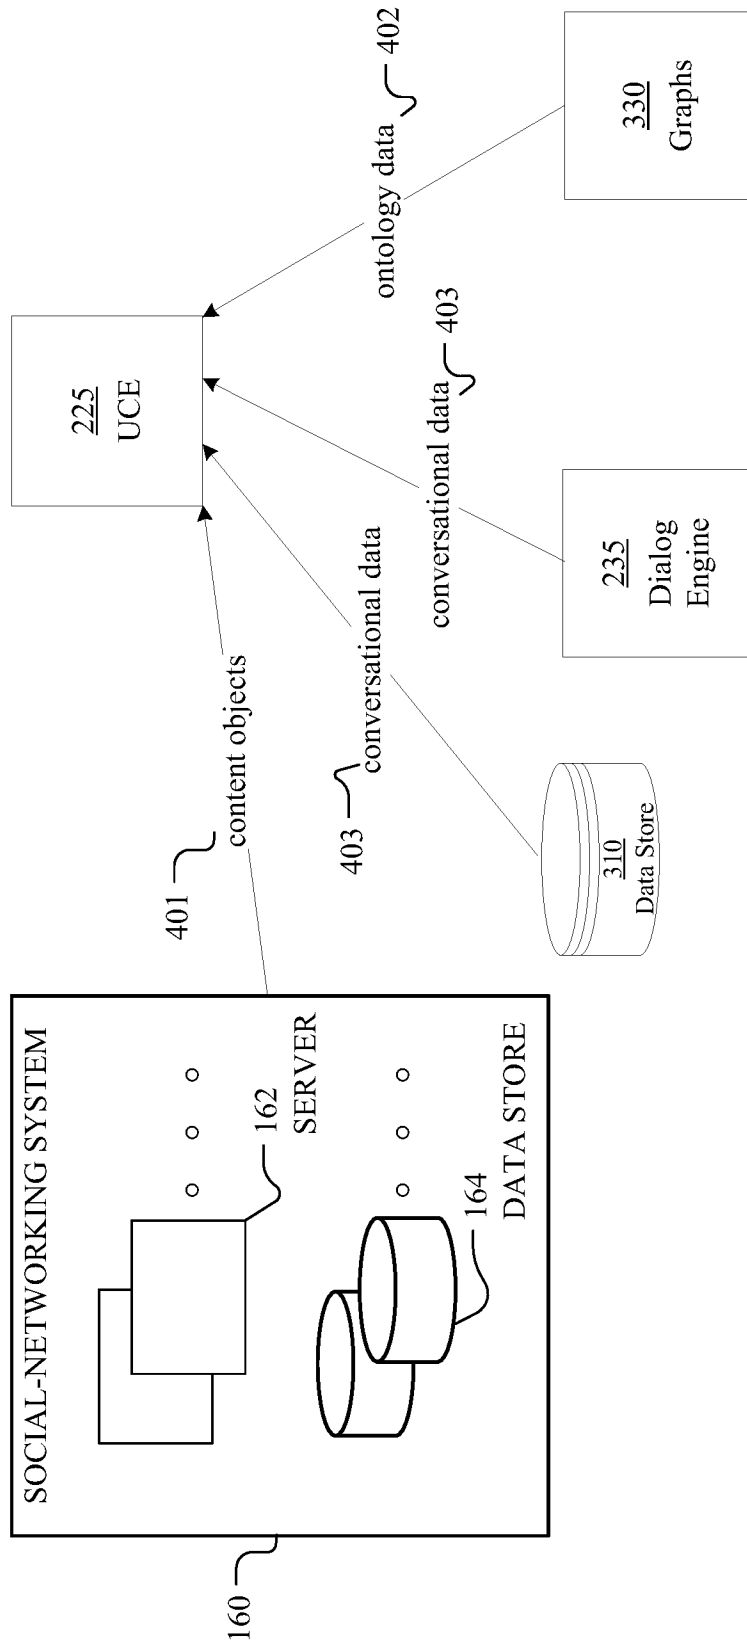


FIG. 4

500

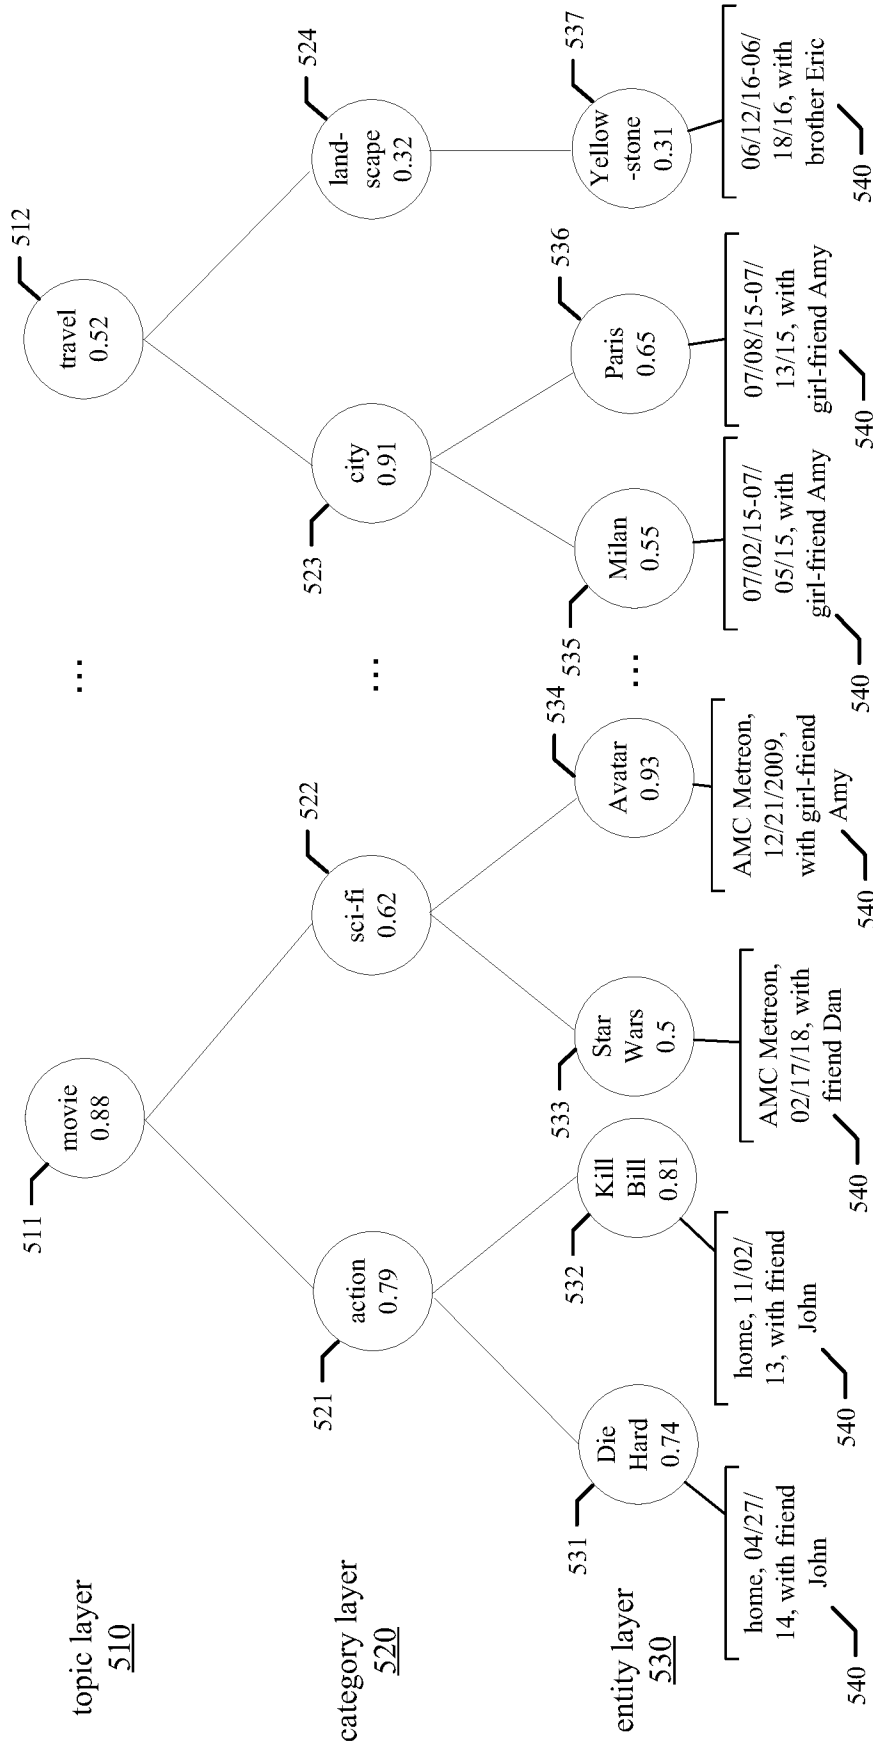
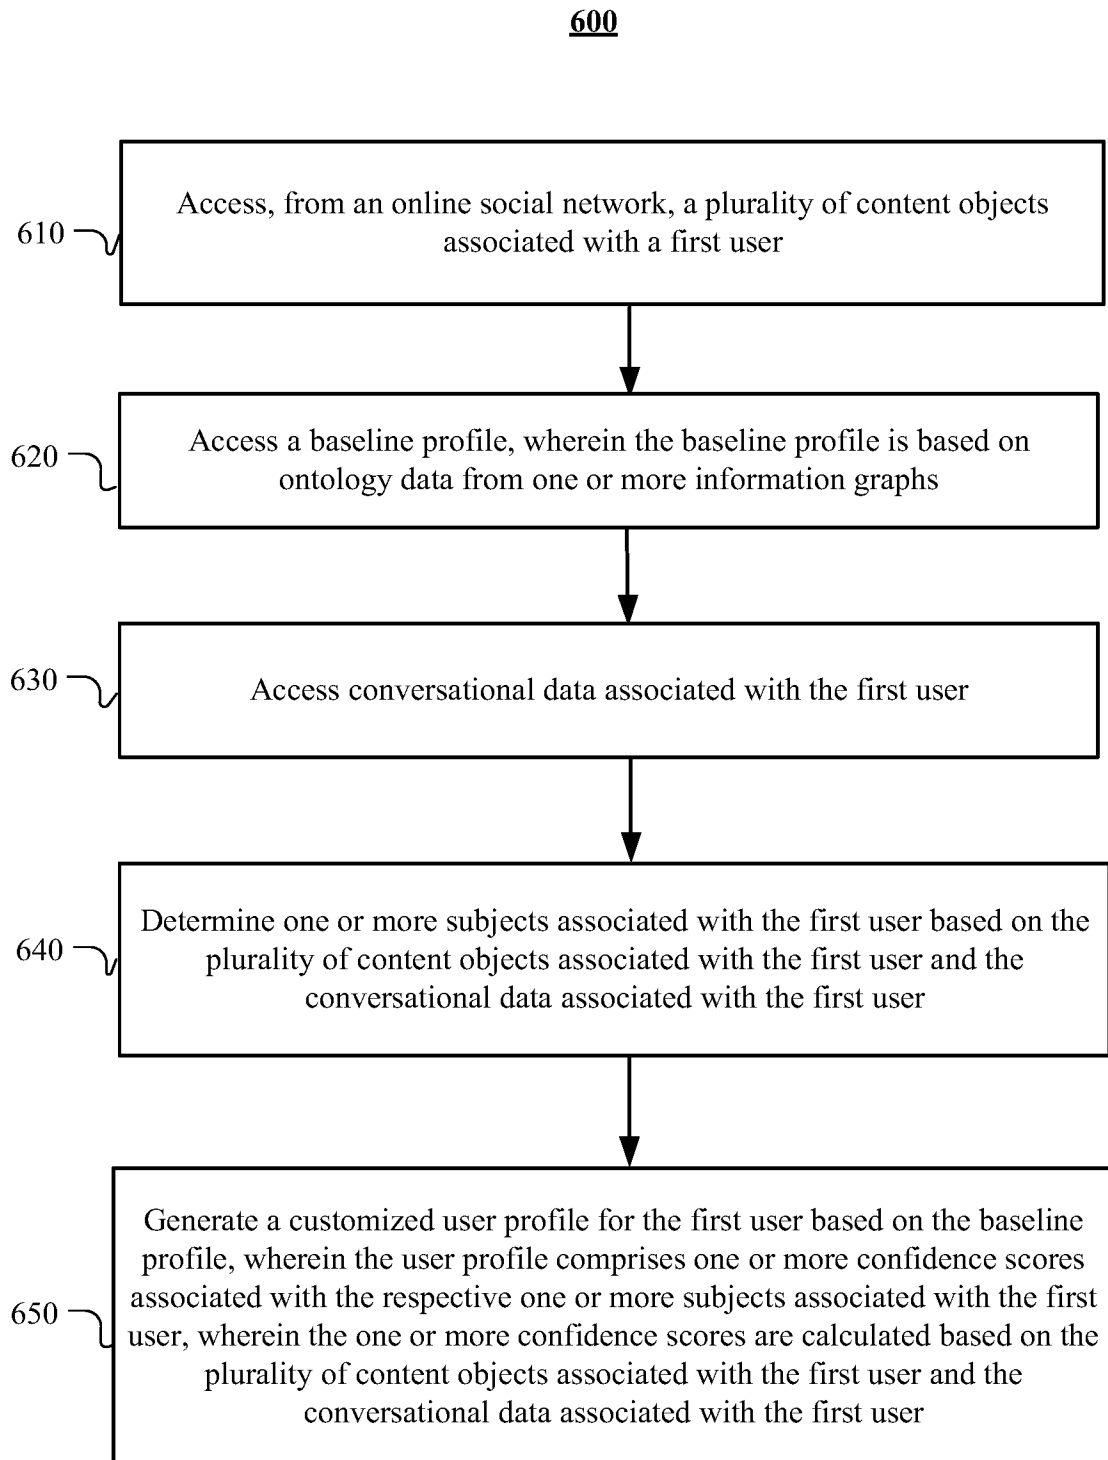


FIG. 5

**FIG. 6**

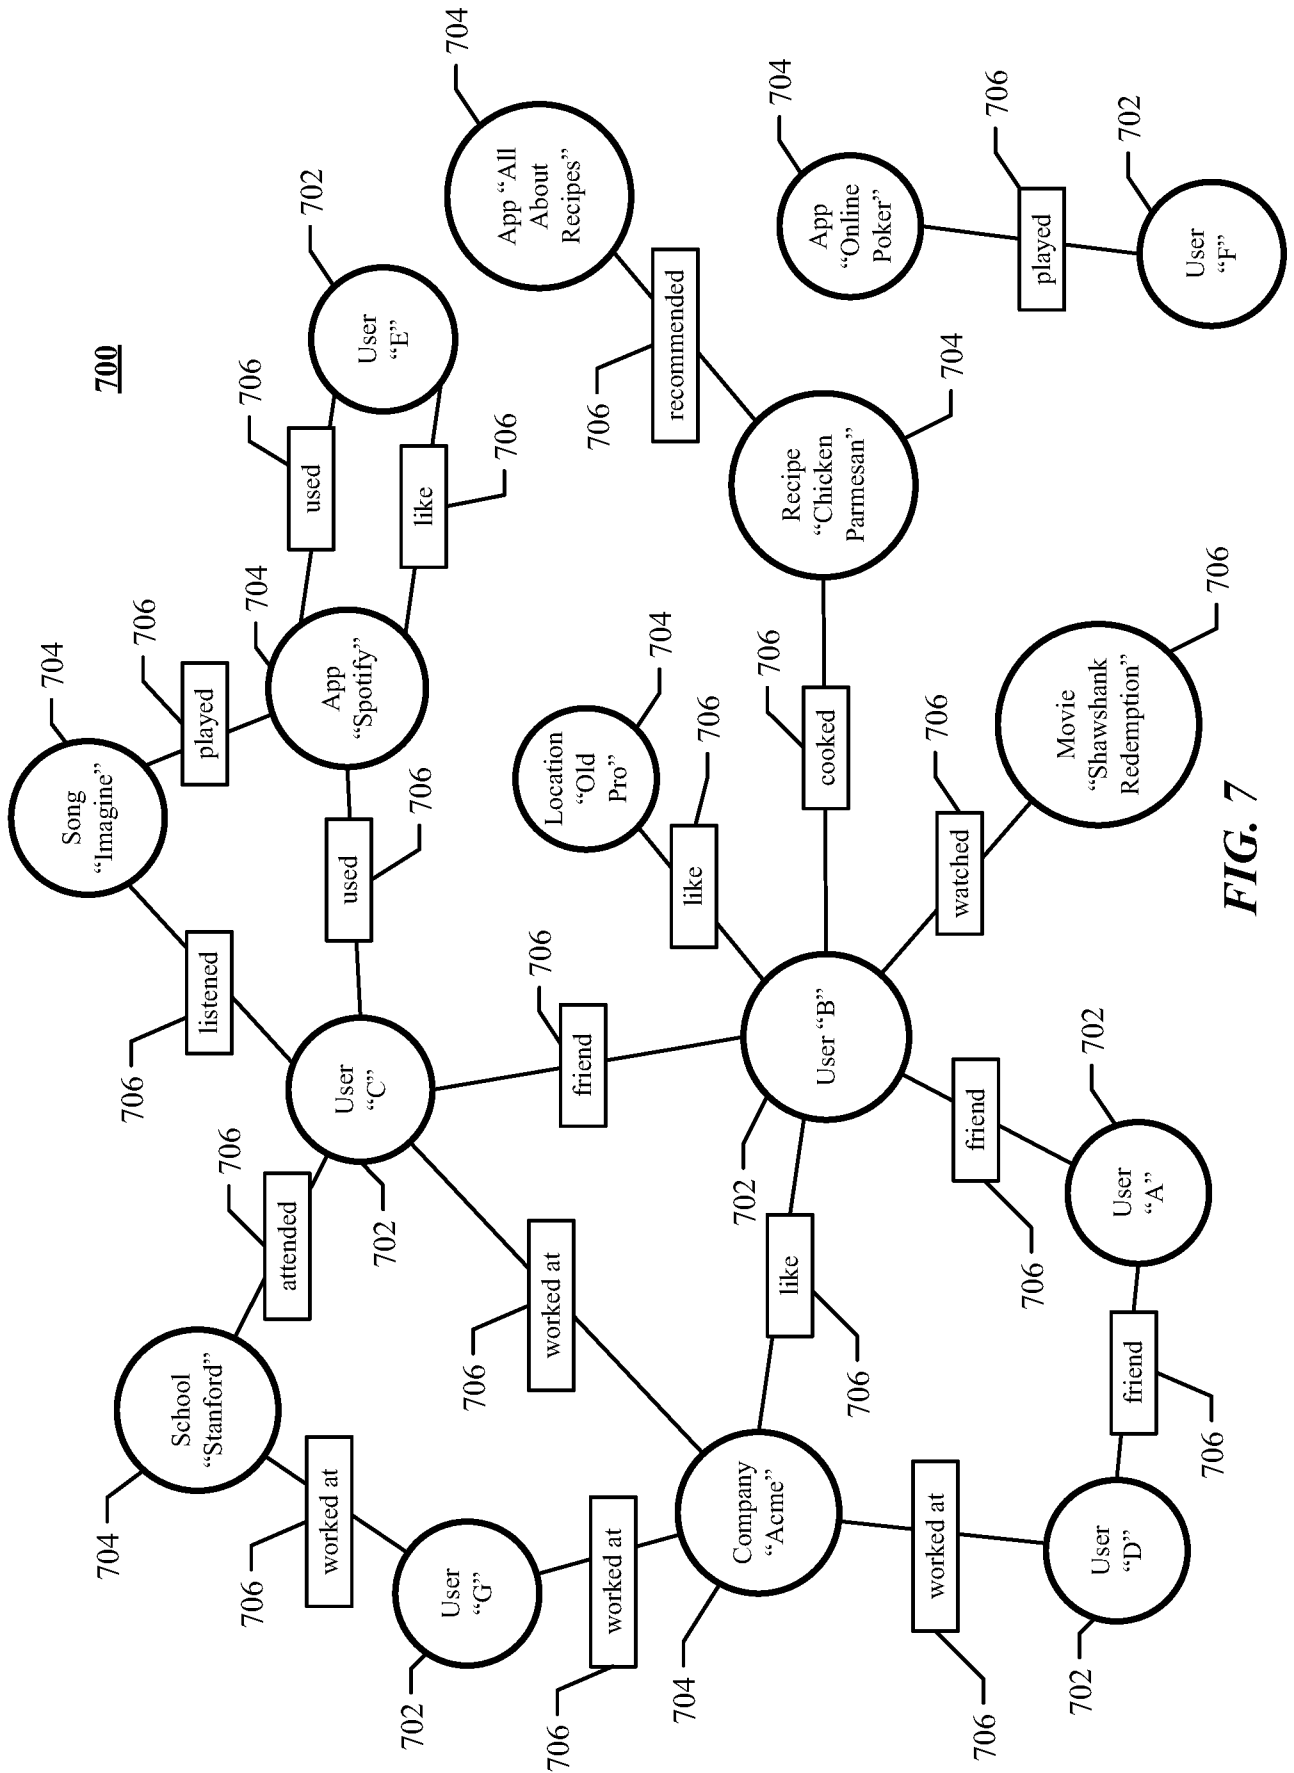


FIG. 7




FIG. 8

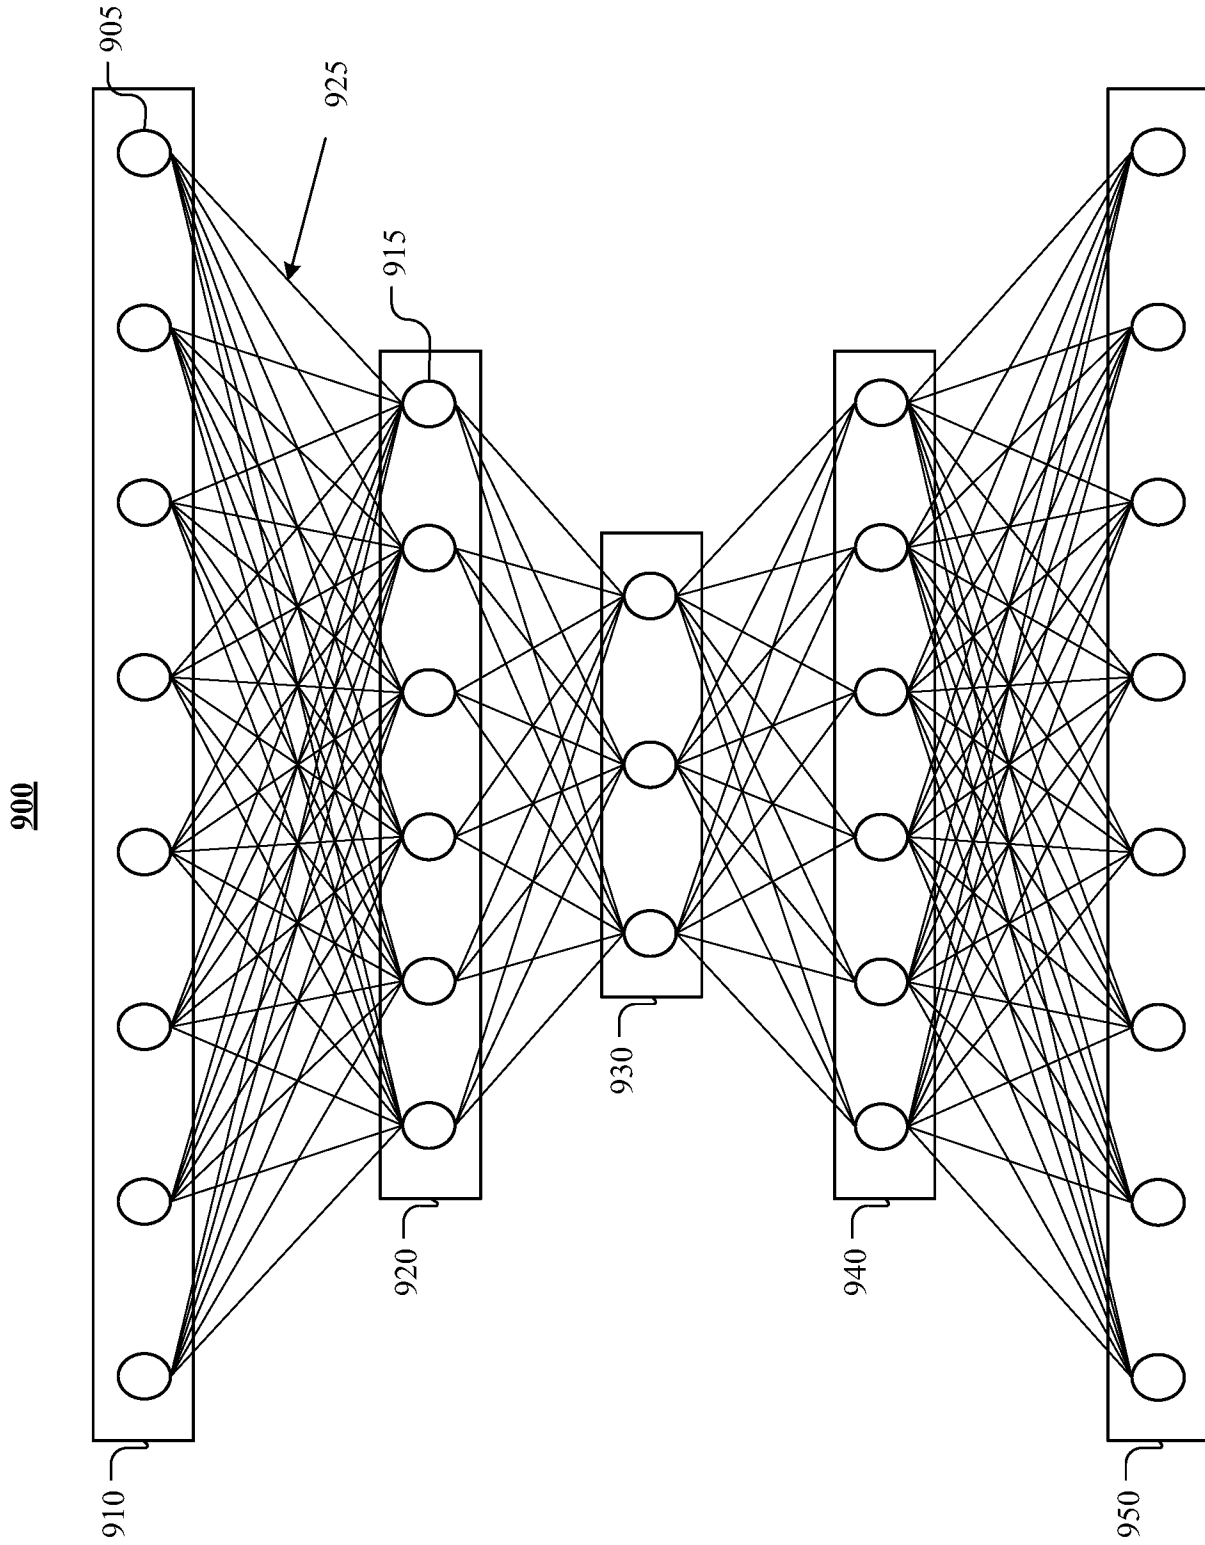


FIG. 9

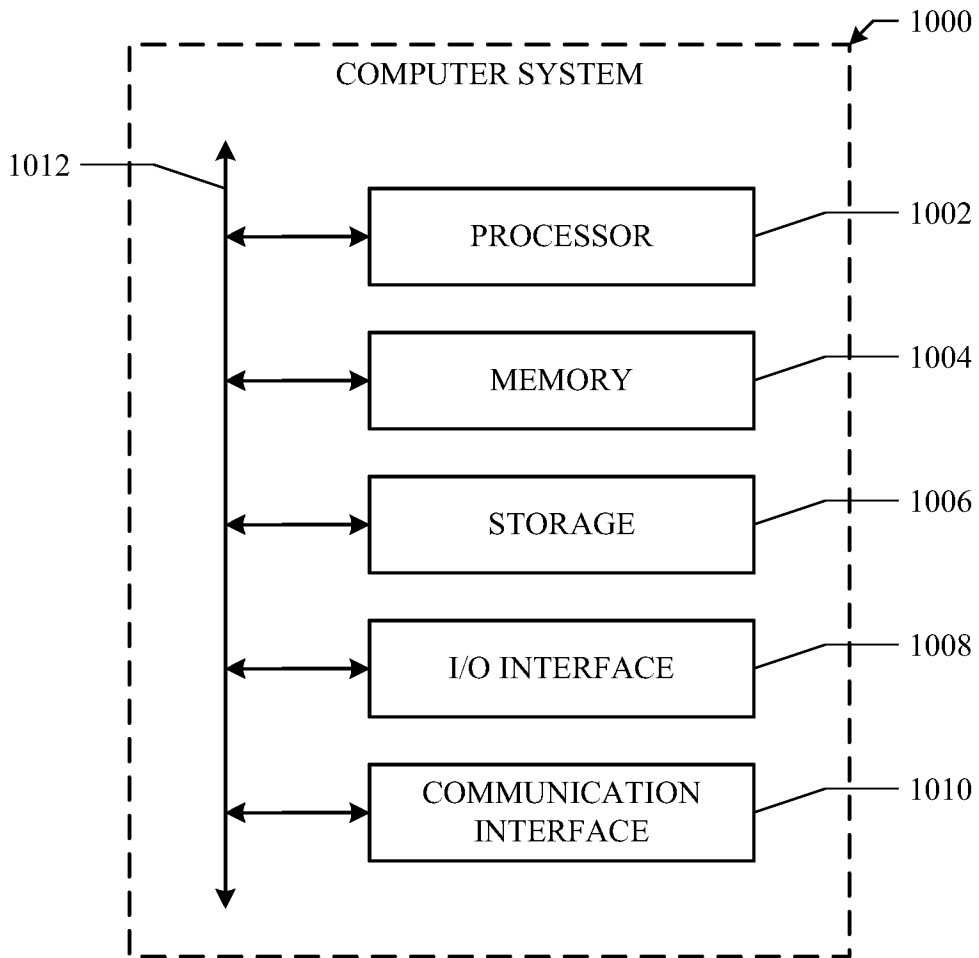


FIG. 10

A. CLASSIFICATION OF SUBJECT MATTER**G06Q 50/00(2006.01)i, G06N 99/00(2010.01)i, G06F 16/00(2019.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06Q 50/00; G06F 17/30; G06N 99/00; G06Q 30/02; H04L 12/58; H04L 29/08; G06F 16/00Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: content, social network, graph, profile, confidence score, conversation, subject**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2015-0006286 A1 (LINKEDIN CORPORATION) 01 January 2015 See paragraphs [0020]-[0022], [0025]-[0026], [0035], [0038]-[0039], [0050]-[0053], [0058], claims 1-2, 7 and figures 1-3.	1-22
Y	US 9824321 B2 (INFOSYS LIMITED) 21 November 2017 See column 1, lines 39-45, column 4, line 47-column 5, line 8, column 6, lines 5-34, claims 1, 16 and figures 5, 8.	1-22
Y	US 2014-0108307 A1 (SRIRAMAN KANDHADAI RAGHUNATHAN et al.) 17 April 2014 See paragraphs [0019]-[0022] and claims 1-9.	8-11, 17
A	US 2015-0032504 A1 (ANBAZHAGAN ELANGO et al.) 29 January 2015 See claims 1-2, 4-6 and figures 1-4.	1-22
A	US 2017-0235726 A1 (FUJITSU LIMITED) 17 August 2017 See claims 1-5 and figures 1-2.	1-22

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

18 January 2019 (18.01.2019)

Date of mailing of the international search report

18 January 2019 (18.01.2019)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

LEE, Myung Jin

Telephone No. +82-42-481-8474



Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:

2. Claims Nos.: 29,32,33,39,44
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
Claims 29, 32, 33, 39, 44 are unclear, because they refer to multiple dependent claims 28, 31, 32, 38, 43 which do not comply with PCT Rule 6.4(a).

3. Claims Nos.: 23-28,30,31,34-38,40-43
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of any additional fees.
3. As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:

4. No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

- Remark on Protest**
- The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2018/039268

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2015-0006286 A1	01/01/2015	US 2014-0244612 A1 US 2015-0006242 A1 US 2015-0006294 A1 US 2015-0006295 A1	28/08/2014 01/01/2015 01/01/2015 01/01/2015
US 9824321 B2	21/11/2017	US 2015-0088593 A1	26/03/2015
US 2014-0108307 A1	17/04/2014	None	
US 2015-0032504 A1	29/01/2015	CN 104081392 A EP 2842057 A1 WO 2013-160904 A1	01/10/2014 04/03/2015 31/10/2013
US 2017-0235726 A1	17/08/2017	JP 2017-142796 A	17/08/2017