



(51) **International Patent Classification:**  
G10L 15/22 (2006.01)

(21) **International Application Number:**  
PCT/US2014/029562

(22) **International Filing Date:**  
14 March 2014 (14.03.2014)

(25) **Filing Language:** English

(26) **Publication Language:** English

(30) **Priority Data:**  
61/792,078 15 March 2013 (15.03.2013) US

(71) **Applicant:** APPLE INC. [US/US]; 1 Infinite Loop, Cupertino, CA 95014 (US).

(72) **Inventors:** PITSCHHEL, Donald, W.; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). CHEYER, Adam, John; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). BRIGHAM, Christopher, Dean; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). GRUBER, Thomas, R.; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US).

(74) **Agents:** EIDE, Christopher B et al; Morrison & Foerster LLP, 755 Page Mill Road, Palo Alto, CA 94304-1018 (US).

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

[Continued on nextpage]

(54) **Title:** TRAINING AN AT LEAST PARTIAL VOICE COMMAND SYSTEM

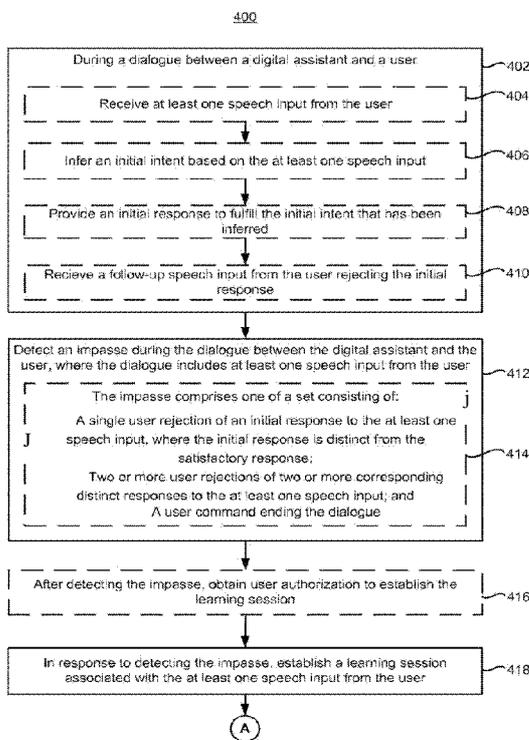


Figure 4A

(57) **Abstract:** An electronic device with one or more processors and memory includes a procedure for training a digital assistant. In some embodiments, the device detects an impasse in a dialogue between the digital assistant and a user including a speech input. During a learning session, the device utilizes a subsequent clarification input from the user to adjust intent inference or task execution associated with the speech input to produce a satisfactory response. In some embodiments, the device identifies a pattern of success or failure associated with an aspect previously used to complete a task and generates a hypothesis regarding a parameter used in speech recognition, intent inference or task execution as a cause for the pattern. Then, the device tests the hypothesis by altering the parameter for a subsequent completion of the task and adopts or rejects the hypothesis based on feedback information collected from the subsequent completion.

WO2014/144949 A2

TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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).

**(84) Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU,

**Published:**

— *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

## TRAINING AN AT LEAST PARTIAL VOICE COMMAND SYSTEM

### CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This Application claims the benefit of U.S. Provisional Application No. 61/792,078, filed on March 15, 2013, entitled TRAINING AN AT LEAST PARTIAL VOICE COMMAND SYSTEM, which is hereby incorporated by reference in its entirety for all purposes.

### TECHNICAL FIELD

[0002] The disclosed embodiments relate generally to digital assistant systems, and more specifically, to training a digital assistant system.

### BACKGROUND

[0003] Just like human personal assistants, digital assistant systems can perform requested tasks and provide requested advice, information, or services. A digital assistant system's ability to fulfill a user's request is dependent on the digital assistant system's correct comprehension of the request or instructions. Recent advances in natural language processing have enabled users to interact with digital assistant systems using natural language, in spoken or textual forms. Such digital assistant systems can interpret the user's input to infer the user's intent, translate the inferred intent into actionable tasks and parameters, execute operations or deploy services to perform the tasks, and produce output that is intelligible to the user. Ideally, the output produced by a digital assistant system should fulfill the user's intent expressed during the natural language interaction between the user and the digital assistant system.

[0004] The ability of a digital assistant system to produce satisfactory responses to user requests depends on the natural language processing, knowledge base, and artificial intelligence available to the digital assistant system. A well-designed training procedure for the digital assistant system can improve a user's experience in interacting with the system and promote the user's confidence in the system's services and capabilities.

## SUMMARY

[0005] The embodiments disclosed herein provide methods, systems, non-transitory computer readable storage medium and user interfaces for training a digital assistant so as to more regularly provide satisfactory responses to a user's requests.

[0006] Accordingly, some embodiments provide a method for operating a digital assistant, the method including, at a device including one or more processors and memory storing one or more programs: detecting an impasse during a dialogue between the digital assistant and a user, where the dialogue includes at least one speech input from the user; and in response to detecting the impasse, establishing a learning session associated with the at least one speech input. During the learning session, the method includes: receiving one or more subsequent clarification inputs from the user; based at least in part on the one or more subsequent clarification inputs, adjusting at least one of intent inference and task execution associated with the at least one speech input to produce a satisfactory response to the at least one speech input; and associating the satisfactory response with the at least one speech input for processing future occurrences of the at least one speech input.

[0007] In some embodiments, another method for training a digital assistant is performed at an electronic device including one or more processors and memory storing instructions for execution by the one or more processors. During a dialogue between the digital assistant and a user, method includes: receiving an initial speech input from a user; inferring an initial intent based on the initial speech input; providing an initial response to fulfill the initial intent that has been inferred; and receiving a follow-up speech input rejecting the initial response. Upon receiving the follow-up speech input rejecting the initial response, the method includes establishing a learning session associated with the initial speech input. During the learning session, the method includes: adjusting at least one of intent inference and task execution associated with the initial speech input to produce a satisfactory response to the initial speech input; and associating the satisfactory response with the initial speech input for processing future occurrences of the initial speech input.

[0008] In some embodiments, a further method for training a digital assistant is performed at an electronic device including one or more processors and memory storing instructions for execution by the one or more processors. The method includes: obtaining feedback information associated with one or more previous completions of a task; and identifying a pattern of success or failure associated with an aspect of speech recognition, intent inference

or task execution previously used to complete the task. The method further includes: generating a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution as a cause for the pattern of success or failure; identifying one or more subsequent requests for completion of the task; testing the hypothesis by altering the parameter used in the at least one of speech recognition, intent inference and task execution for subsequent completions of the task; and adopting or rejecting the hypothesis based on feedback information collected from the subsequent completions of the task.

[0009] In another aspect, an electronic device includes one or more processors and memory storing one or more programs for execution by the one or more processors, where the one or more programs include instructions that when executed by the one or more processors cause the electronic device to perform any of the aforementioned methods.

[0010] In yet another aspect, a non-transitory computer readable medium stores one or more programs that when executed by one or more processors of a computer system cause the electronic device to perform any of the aforementioned methods.

[0011] The details of one or more embodiments of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0012] Figure 1 is a block diagram illustrating an environment in which a digital assistant operates in accordance with some embodiments.

[0013] Figure 2 is a block diagram illustrating a digital assistant client system in accordance with some embodiments.

[0014] Figure 3A is a block diagram illustrating a digital assistant system or a server portion thereof in accordance with some embodiments.

[0015] Figure 3B is a block diagram illustrating functions of the digital assistant shown in Figure 3A in accordance with some embodiments.

[0016] Figure 3C is a diagram of a portion of an ontology shown in Figure 3B in accordance with some embodiments.

[0017] Figures 4A-4C are a flow chart for an exemplary process for training a digital assistant in accordance with some embodiments.

[0018] Figure 5 is a block diagram of a training module included in Figure 3B in accordance with some embodiments.

[0019] Figures 6A-6C are a flow chart for an exemplary process for training a digital assistant in accordance with some embodiments.

[0020] Figure 7 is a functional block diagram of an electronic device in accordance with some embodiments.

[0021] Figure 8 is a functional block diagram of an electronic device in accordance with some embodiments.

[0022] Like reference numerals refer to corresponding parts throughout the drawings.

#### DESCRIPTION OF EMBODIMENTS

[0023] Figure 1 is a block diagram of an operating environment 100 of a digital assistant according to some embodiments. The terms "digital assistant," "virtual assistant," "intelligent automated assistant," or "automatic digital assistant," refer to any information processing system that interprets natural language input in spoken and/or textual form to infer user intent, and performs actions based on the inferred user intent. For example, to act on an inferred user intent, the system can perform one or more of the following: identifying a task flow with steps and parameters designed to accomplish the inferred user intent, inputting specific requirements from the inferred user intent into the task flow; executing the task flow by invoking programs, methods, services, APIs, or the like; and generating output responses to the user in an audible (e.g. speech) and/or visual form.

[0024] Specifically, a digital assistant is capable of accepting a user request at least partially in the form of a natural language command, request, statement, narrative, and/or inquiry. Typically, the user request seeks either an informational answer or performance of a task by the digital assistant. A satisfactory response to the user request is either provision of the requested informational answer, performance of the requested task, or a combination of the two. For example, a user may ask the digital assistant a question, such as "Where am I right now?" Based on the user's current location, the digital assistant may answer, "You are in Central Park." The user may also request the performance of a task, for example, "Please remind me to call mom at 4pm today." In response, the digital assistant may acknowledge the

request and then creates an appropriate reminder item in the user's electronic schedule. During performance of a requested task, the digital assistant sometimes interacts with the user in a continuous dialogue involving multiple exchanges of information over an extended period of time. There are numerous other ways of interacting with a digital assistant to request information or performance of various tasks. In addition to providing verbal responses and taking programmed actions, the digital assistant also provides responses in other visual or audio forms (e.g., as text, alerts, music, videos, animations, etc.).

**[0025]** An example of a digital assistant is described in Applicant's U.S. Utility Application Serial No. 12/987,982 for "Intelligent Automated Assistant," filed January 10, 2011, the entire disclosure of which is incorporated herein by reference.

**[0026]** As shown in Figure 1, in some embodiments, a digital assistant is implemented according to a client-server model. The digital assistant includes a client-side portion 102a, 102b (hereafter "DA-client 102") executed on a user device 104a, 104b, and a server-side portion 106 (hereafter "DA-server 106") executed on a server system 108. The DA-client 102 communicates with the DA-server 106 through one or more networks 110. The DA-client 102 provides client-side functionalities such as user-facing input and output processing and communications with the DA-server 106. The DA server 106 provides server-side functionalities for any number of DA-clients 102 each residing on a respective user device 104.

**[0027]** In some embodiments, the DA-server 106 includes a client-facing I/O interface 112, one or more processing modules 114, data and models 116, and an interface 170 to external services 118. The client-facing interface 170 facilitates the client-facing input and output processing for the digital assistant server 106. The one or more processing modules 114 utilize the data and models 116 to determine the user's intent based on natural language input and perform task execution based on inferred user intent. In some embodiments, data and models 116 stores a plurality of ontologies (e.g., a global ontology, a regional ontology, a cultural ontology, a national ontology, a state-wide ontology, a city-wide ontology, etc.) maintained by DA-server 106 distinct from a respective user's ontology (e.g., ontology 360). The functions of an ontology are described in more detail below with respect to Figure 3B. In some embodiments, the DA-server 106 communicates with external services 120 through the network(s) 110 for task completion or information acquisition. The interface 170 to external services 118 facilitates such communications.

[0028] Examples of the user device 104 include, but are not limited to, a handheld computer, a personal digital assistant (PDA), a tablet computer, a laptop computer, a desktop computer, a cellular telephone, a smart phone, an enhanced general packet radio service (EGPRS) mobile phone, a media player, a navigation device, a game console, a television, a remote control, or a combination of any two or more of these data processing devices or other data processing devices. More details on the user device 104 are provided in reference to an exemplary user device 104 shown in Figure 2.

[0029] Examples of the communication network(s) 110 include local area networks ("LAN") and wide area networks ("WAN"), e.g., the Internet. The communication network(s) 110 are, optionally, implemented using any known network protocol, including various wired or wireless protocols, such as e.g., Ethernet, Universal Serial Bus (USB), **FIPvEWIRE**, Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Wi-Fi, voice over Internet Protocol (VoIP), Wi-MAX, or any other suitable communication protocol.

[0030] The server system 108 is implemented on one or more standalone data processing apparatus or a distributed network of computers. In some embodiments, the server system 108 also employs various virtual devices and/or services of third party service providers (e.g., third-party cloud service providers) to provide the underlying computing resources and/or infrastructure resources of the server system 108.

[0031] Although the digital assistant shown in Figure 1 includes both a client-side portion (e.g., the DA-client 102) and a server-side portion (e.g., the DA-server 106), in some embodiments, the functions of a digital assistant is implemented as a standalone application installed on a user device. In addition, the divisions of functionalities between the client and server portions of the digital assistant can vary in different embodiments. For example, in some embodiments, the DA-client is a thin-client that provides only user-facing input and output processing functions, and delegates all other functionalities of the digital assistant to a backend server.

[0032] As described later in this specification, the digital assistant can implement a crowd sourcing functionality. The crowd sourcing functionality allows the digital assistant to gather information from other DA-clients or third party information sources (so-called "crowd-

sourcing information sources" or "CS information sources"), and use the crowd sourced information to facilitate request fulfillment, and in some cases, intent inference.

**[0033]** Figure 2 is a block diagram of a user-device 104 in accordance with some embodiments. The user device 104 includes a memory interface 202, one or more processors 204, and a peripherals interface 206. The various components in the user device 104 are coupled by one or more communication buses or signal lines. The user device 104 includes various sensors, subsystems, and peripheral devices that are coupled to the peripherals interface 206. The sensors, subsystems, and peripheral devices gather information and/or facilitate various functionalities of the user device 104.

**[0034]** For example, a motion sensor 210, a light sensor 212, and a proximity sensor 214 are coupled to the peripherals interface 206 to facilitate orientation, light, and proximity sensing functions. One or more other sensors 216, such as a positioning system (e.g., a GPS receiver), a temperature sensor, a biometric sensor, a gyroscope, a compass, an accelerometer, and the like, are also connected to the peripherals interface 206, to facilitate related functionalities.

**[0035]** In some embodiments, a camera subsystem 220 and an optical sensor 222 are utilized to facilitate camera functions, such as taking photographs and recording video clips. Communication functions are facilitated through one or more wired and/or wireless communication subsystems 224, which can include various communication ports, radio frequency receivers and transmitters, and/or optical (e.g., infrared) receivers and transmitters. An audio subsystem 226 is coupled to speakers 228 and a microphone 230 to facilitate voice-enabled functions, such as voice recognition, voice replication, digital recording, and telephony functions.

**[0036]** In some embodiments, an I/O subsystem 240 is also coupled to the peripheral interface 206. The I/O subsystem 240 includes a touch screen controller 242 and/or other input controller(s) 244. The touch-screen controller 242 is coupled to a touch screen 246. The touch screen 246 and the touch screen controller 242 can, for example, detect contact and movement or break thereof using any of a plurality of touch sensitivity technologies, such as capacitive, resistive, infrared, surface acoustic wave technologies, proximity sensor arrays, and the like. The other input controller(s) 244 can be coupled to other input/control devices 248, such as one or more buttons, rocker switches, thumb-wheel, infrared port, USB port, and/or a pointer device such as a stylus.

[0037] In some embodiments, the memory interface 202 is coupled to memory 250. The memory 250 can include high-speed random access memory and/or non-volatile memory, such as one or more magnetic disk storage devices, one or more optical storage devices, and/or flash memory (e.g., NAND, NOR).

[0038] In some embodiments, the memory 250 stores an operating system 252, a communication module 254, a graphical user interface module 256, a sensor processing module 258, a phone module 260, and applications 262. The operating system 252 includes instructions for handling basic system services and for performing hardware dependent tasks. The communication module 254 facilitates communicating with one or more additional devices, one or more computers and/or one or more servers. The graphical user interface module 256 facilitates graphic user interface processing. The sensor processing module 258 facilitates sensor-related processing and functions. The phone module 260 facilitates phone-related processes and functions. The application module 262 facilitates various functionalities of user applications, such as electronic-messaging, web browsing, media processing, Navigation, imaging and/or other processes and functions.

[0039] As described in this specification, the memory 250 also stores client-side digital assistant instructions (e.g., in a digital assistant client module 264) and various user data 266 (e.g., user-specific vocabulary data, preference data, and/or other data such as the user's electronic address book, to-do lists, shopping lists, etc.) to provide the client-side functionalities of the digital assistant.

[0040] In various embodiments, the digital assistant client module 264 is capable of accepting voice input (e.g., speech input), text input, touch input, and/or gestural input through various user interfaces (e.g., the I/O subsystem 244) of the user device 104. The digital assistant client module 264 is also capable of providing output in audio (e.g., speech output), visual, and/or tactile forms. For example, output can be provided as voice, sound, alerts, text messages, menus, graphics, videos, animations, vibrations, and/or combinations of two or more of the above. During operation, the digital assistant client module 264 communicates with the digital assistant server using the communication subsystems 224.

[0041] In some embodiments, the digital assistant client module 264 utilizes the various sensors, subsystems and peripheral devices to gather additional information from the surrounding environment of the user device 104 to establish a context associated with a user, the current user interaction, and/or the current user input. In some embodiments, the digital

assistant client module 264 provides the context information or a subset thereof with the user input to the digital assistant server to help infer the user's intent. In some embodiments, the digital assistant also uses the context information to determine how to prepare and delivery outputs to the user.

**[0042]** In some embodiments, the context information that accompanies the user input includes sensor information, e.g., lighting, ambient noise, ambient temperature, images or videos of the surrounding environment, etc. In some embodiments, the context information also includes the physical state of the user device 104 (e.g., device orientation, device location, device temperature, power level, speed, acceleration, motion patterns, cellular signals strength, etc.). In some embodiments, information related to the software state of the user device 104 (e.g., running processes, installed programs, past and present network activities, background services, error logs, resources usage, etc.) is provided to the digital assistant server as context information associated with a user input.

**[0043]** In some embodiments, the digital assistant client module 264 selectively provides information (e.g., user data 266) stored on the user device 104 in response to requests from the digital assistant server. In some embodiments, the digital assistant client module 264 also elicits additional input from the user via a natural language dialogue or other user interfaces upon request by the digital assistant server 106. The digital assistant client module 264 passes the additional input to the digital assistant server 106 to help the digital assistant server 106 in intent inference and/or fulfillment of the user's intent expressed in the user request.

**[0044]** In various embodiments, the memory 250 includes additional instructions or fewer instructions. Furthermore, various functions of the user device 104 may be implemented in hardware and/or in firmware, including in one or more signal processing and/or application specific integrated circuits.

**[0045]** Figure 3A is a block diagram of an example digital assistant system 300 in accordance with some embodiments. In some embodiments, the digital assistant system 300 is implemented on a standalone computer system. In some embodiments, the digital assistant system 300 is distributed across multiple computers. In some embodiments, some of the modules and functions of the digital assistant are divided into a server portion and a client portion, where the client portion resides on a user device (e.g., the user device 104) and communicates with the server portion (e.g., the server system 108) through one or more networks (e.g., as shown in Figure 1). In some embodiments, the digital assistant system 300

is an embodiment of the server system 108 (and/or the digital assistant server 106) shown in Figure 1. It should be noted that the digital assistant system 300 is only one example of a digital assistant system, and that the digital assistant system 300 may have more or fewer components than shown, may combine two or more components, or may have a different configuration or arrangement of the components. The various components shown in Figure 3A may be implemented in hardware, software instructions for execution by one or more processors, firmware, including one or more signal processing and/or application specific integrated circuits, or a combination of thereof.

**[0046]** The digital assistant system 300 includes memory 302, one or more processors 304, an input/output (I/O) interface 306, and a network communications interface 308. These components communicate with one another over one or more communication buses or signal lines 310.

**[0047]** In some embodiments, the memory 302 includes a non-transitory computer readable storage medium, such as high-speed random access memory and/or a non-volatile storage medium (e.g., one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices).

**[0048]** In some embodiments, the I/O interface 306 couples input/output devices 316 of the digital assistant system 300, such as displays, keyboards, touch screens, and microphones, to the user interface module 322. The I/O interface 306, in conjunction with the user interface module 322, receives user inputs (e.g., voice input, keyboard inputs, touch inputs, etc.) and processes them accordingly. In some embodiments, the digital assistant system 300 includes any of the components and I/O and communication interfaces described with respect to the user device 104 in Figure 2 (e.g., when the digital assistant is implemented on a standalone user device.). In some embodiments, the digital assistant system 300 represents the server portion of a digital assistant implementation, and interacts with the user through a client-side portion residing on a user device (e.g., the user device 104 shown in Figure 2).

**[0049]** In some embodiments, the network communications interface 308 includes wired communication port(s) 312 and/or wireless transmission and reception circuitry 314. The wired communication port(s) receive and send communication signals via one or more wired interfaces, e.g., Ethernet, Universal Serial Bus (USB), FIREWIRE, etc. The wireless circuitry 314 receives and sends RF signals and/or optical signals from/to communications networks and other communications devices. The wireless communications, optionally, use

any of a plurality of communications standards, protocols and technologies, such as GSM, EDGE, CDMA, TDMA, Bluetooth, Wi-Fi, VoIP, Wi-MAX, or any other suitable communication protocol. The network communications interface 308 enables communication between the digital assistant system 300 with networks, such as the Internet, an intranet and/or a wireless network, such as a cellular telephone network, a wireless local area network (LAN) and/or a metropolitan area network (MAN), and other devices.

**[0050]** In some embodiments, memory 302, or the computer readable storage media of memory 302, stores programs, modules, instructions, and data structures including all or a subset of: an operating system 318, a communications module 320, a user interface module 322, one or more applications 324, and a digital assistant module 326. The one or more processors 304 execute these programs, modules, and instructions, and reads/writes from/to the data structures.

**[0051]** The operating system 318 (e.g., Darwin, RTXC, LINUX, UNIX, OS X, WINDOWS, or an embedded operating system such as VxWorks) includes various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communications between various hardware, firmware, and software components.

**[0052]** The communications module 320 facilitates communications over the network communications interface 308 between the digital assistant system 300 and other devices. For example, the communication module 320, optionally, communicates with the communication interface 254 of the device 104 shown in Figure 2. The communications module 320 also includes various components for handling data received by the wireless circuitry 314 and/or wired communications port 312.

**[0053]** The user interface module 322 receives commands and/or inputs from a user via the I/O interface 306 (e.g., from a keyboard, touch screen, pointing device, controller, and/or microphone), and generates user interface objects on a display. The user interface module 322 also prepares and delivers outputs (e.g., speech, sound, animation, text, icons, vibrations, haptic feedback, and light, etc.) to the user via the I/O interface 306 (e.g., through displays, audio channels, speakers, and touch-pads, etc.).

**[0054]** The applications 324 include programs and/or modules that are configured to be executed by the one or more processors 304. For example, if the digital assistant system is implemented on a standalone user device, the applications 324, optionally, include user

applications, such as games, a calendar application, a navigation application, or an email application. If the digital assistant system 300 is implemented on a server farm, the applications 324, optionally, include resource management applications, diagnostic applications, or scheduling applications, for example.

**[0055]** The memory 302 also stores the digital assistant module (or the server portion of a digital assistant) 326. In some embodiments, the digital assistant module 326 includes the following sub-modules, or a subset or superset thereof: an input/output processing module 328, a speech-to-text (STT) processing module 330, a natural language processing module 332, a dialogue flow processing module 334, a task flow processing module 336, a service processing module 338, a training module 340, and a crowd sourcing module 342. Each of these modules has access to one or more of the following data and models of the digital assistant 326, or a subset or superset thereof: ontology 360, vocabulary index 344, user data 348, categorization module 349, disambiguation module 350, task flow models 354, service models 356, CS knowledge base 358, and user log 370.

**[0056]** In some embodiments, using the processing modules, data, and models implemented in the digital assistant module 326, the digital assistant system 300 performs at least some of the following: identifying a user's intent expressed in a natural language input received from the user; actively eliciting and obtaining information needed to fully infer the user's intent (e.g., by disambiguating words, names, intentions, etc.); determining the task flow for fulfilling the inferred intent; and executing the task flow to fulfill the inferred intent. In some embodiments, the digital assistant also takes appropriate actions when a satisfactory response was not or could not be provided to the user for various reasons.

**[0057]** In some embodiments, as shown in Figure 38, the I/O processing module 328 interacts with the user through the I/O devices 316 in Figure 3A or with a user device (e.g., a user device 104 in Figure 1) through the network communications interface 308 in Figure 3A to obtain user input (e.g., a speech input) and to provide responses (e.g., as speech outputs) to the user input. The I/O processing module 328, optionally, obtains context information associated with the user input from the user device, along with or shortly after the receipt of the user input. The context information includes user-specific data, vocabulary, and/or preferences relevant to the user input. In some embodiments, the context information also includes software and hardware states of the device (e.g., the user device 104 in Figure 1) at the time the user request is received, and/or information related to the surrounding environment of the user at the time that the user request was received. In some embodiments,

the I/O processing module 328 also sends follow-up questions to, and receives answers from, the user regarding the user request. When a user request is received by the I/O processing module 328 and the user request contains a speech input, the I/O processing module 328 forwards the speech input to the speech-to-text (STT) processing module 330 for speech-to-text conversion.

**[0058]** The speech-to-text processing module 330 receives speech input (e.g., a user utterance captured in a voice recording) through the I/O processing module 328. In some embodiments, the speech-to-text processing module 330 uses various acoustic and language models to recognize the speech input as a sequence of phonemes, and ultimately, a sequence of words or tokens written in one or more languages. The speech-to-text processing module 330 can be implemented using any suitable speech recognition techniques, acoustic models, and language models, such as Hidden Markov Models, Dynamic Time Warping (DTW) based speech recognition, and other statistical and/or analytical techniques. In some embodiments, the speech-to-text processing can be performed at least partially by a third party service or on the user's device. Once the speech-to-text processing module 330 obtains the result of the speech-to-text processing (e.g., a sequence of words or tokens) it passes the result to the natural language processing module 332 for intent inference.

**[0059]** More details on the speech-to-text processing are described in U.S. Utility Application Serial No. 13/236,942 for "Consolidating Speech Recognition Results," filed on September 20, 2011, the entire disclosure of which is incorporated herein by reference.

**[0060]** The natural language processing module 332 ("natural language processor") of the digital assistant takes the sequence of words or tokens ("token sequence") generated by the speech-to-text processing module 330, and attempts to associate the token sequence with one or more "actionable intents" recognized by the digital assistant. An "actionable intent" represents a task that can be performed by the digital assistant, and has an associated task flow implemented in the task flow models 354. The associated task flow is a series of programmed actions and steps that the digital assistant takes in order to perform the task. The scope of a digital assistant's capabilities is dependent on the number and variety of task flows that have been implemented and stored in the task flow models 354, or in other words, on the number and variety of "actionable intents" that the digital assistant recognizes. The effectiveness of the digital assistant, however, is also dependent on the assistant's ability to infer the correct "actionable intent(s)" from the user request expressed in natural language.

[0061] In some embodiments, in addition to the sequence of words or tokens obtained from the speech-to-text processing module 330, the natural language processor 332 also receives context information associated with the user request (e.g., from the I/O processing module 328). The natural language processor 332, optionally, uses the context information to clarify, supplement, and/or further define the information contained in the token sequence received from the speech-to-text processing module 330. The context information includes, for example, user preferences, hardware and/or software states of the user device, sensor information collected before, during, or shortly after the user request, prior interactions (e.g., dialogue) between the digital assistant and the user, and the like.

[0062] In some embodiments, the natural language processing is based on ontology 360. The ontology 360 is a hierarchical structure containing many nodes, each node representing either an "actionable intent" or a "property" relevant to one or more of the "actionable intents" or other "properties." As noted above, an "actionable intent" represents a task that the digital assistant is capable of performing (i.e., it is "actionable" or can be acted on). A "property" represents a parameter associated with an actionable intent or a sub-aspect of another property. A linkage between an actionable intent node and a property node in the ontology 360 defines how a parameter represented by the property node pertains to the task represented by the actionable intent node.

[0063] In some embodiments, the ontology 360 is made up of actionable intent nodes and property nodes. Within the ontology 360, each actionable intent node is linked to one or more property nodes either directly or through one or more intermediate property nodes. Similarly, each property node is linked to one or more actionable intent nodes either directly or through one or more intermediate property nodes. For example, as shown in Figure 3C, the ontology 360 optionally includes a "restaurant reservation" node - an actionable intent node. Property nodes "restaurant," "date/time" (for the reservation), and "party size" are each directly linked to the actionable intent node (e.g., the "restaurant reservation" node). In addition, property nodes "cuisine," "price range," "phone number," and "location" are sub-nodes of the property node "restaurant," and are each linked to the "restaurant reservation" node (e.g., the actionable intent node) through the intermediate property node "restaurant." For another example, as shown in Figure 3C, the ontology 360 may also include a "set reminder" node (e.g., another actionable intent node). Property nodes "date/time" (for the setting the reminder) and "subject" (for the reminder) are each linked to the "set reminder" node. Since the property "date/time" is relevant to both the task of making a

restaurant reservation and the task of setting a reminder, the property node "date/time" is linked to both the "restaurant reservation" node and the "set reminder" node in the ontology 360.

**[0064]** An actionable intent node, along with its linked property nodes, is sometimes described as a "domain." In the present discussion, each domain is associated with a respective actionable intent, and refers to the group of nodes (and the relationships therebetween) associated with the particular actionable intent. For example, the ontology 360 shown in Figure 3C includes an example of a restaurant reservation domain 362 and an example of a reminder domain 364 within the ontology 360. The restaurant reservation domain includes the actionable intent node "restaurant reservation," property nodes "restaurant," "date/time," and "party size," and sub-property nodes "cuisine," "price range," "phone number," and "location." The reminder domain 364 includes the actionable intent node "set reminder," and property nodes "subject" and "date/time." In some embodiments, the ontology 360 is made up of many domains. Each domain optionally shares one or more property nodes with one or more other domains. For example, the "date/time" property node is optionally associated with many different domains (e.g., a scheduling domain, a travel reservation domain, a movie ticket domain, etc.), in addition to the restaurant reservation domain 362 and the reminder domain 364.

**[0065]** While Figure 3C illustrates two example domains within the ontology 360, other domains (or actionable intents) include, for example, "initiate a phone call," "find directions," "schedule a meeting," "send a message," and "provide an answer to a question," and so on. A "send a message" domain is associated with a "send a message" actionable intent node, and optionally further includes property nodes such as "recipient(s)," "message type," and "message body." The property node "recipient" is optionally further defined, for example, by the sub-property nodes such as "recipient name" and "message address."

**[0066]** In some embodiments, the ontology 360 includes all the domains (and hence actionable intents) that the digital assistant is capable of understanding and acting upon. In some embodiments, the ontology 360 is optionally modified, such as by adding or removing entire domains or nodes, or by modifying relationships between the nodes within the ontology 360.

**[0067]** In some embodiments, nodes associated with multiple related actionable intents are optionally clustered under a "super domain" in the ontology 360. For example, a "travels"

super-domain optionally includes a cluster of property nodes and actionable intent nodes related to travels. The actionable intent nodes related to travels optionally includes "airline reservation," "hotel reservation," "car rental," "get directions," "find points of interest," and so on. The actionable intent nodes under the same super domain (e.g., the "travels" super domain) sometimes have many property nodes in common. For example, the actionable intent nodes for "airline reservation," "hotel reservation," "car rental," "get directions," "find points of interest" sometimes share one or more of the property nodes "start location," "destination," "departure date/time," "arrival date/time," and "party size."

**[0068]** In some embodiments, each node in the ontology 360 is associated with a set of words and/or phrases that are relevant to the property or actionable intent represented by the node. The respective set of words and/or phrases associated with each node is the so-called "vocabulary" associated with the node. The respective set of words and/or phrases associated with each node can be stored in the vocabulary index 344 in association with the property or actionable intent represented by the node. For example, returning to Figure 3B, the vocabulary associated with the node for the property of "restaurant" optionally includes words such as "food," "drinks," "cuisine," "hungry," "eat," "pizza," "fast food," "meal," and so on. For another example, the vocabulary associated with the node for the actionable intent of "initiate a phone call" optionally includes words and phrases such as "call," "phone," "dial," "call this number," "make a call to," and so on. The vocabulary index 344, optionally, includes words and phrases in different languages.

**[0069]** The natural language processor 332 receives the token sequence (e.g., a text string) from the speech-to-text processing module 330, and determines what nodes are implicated by the words in the token sequence. In some embodiments, if a word or phrase in the token sequence is found to be associated with one or more nodes in the ontology 360 (via the vocabulary index 344), the word or phrase will "trigger" or "activate" those nodes. Based on the quantity and/or relative importance of the activated nodes, the natural language processor 332 will select one of the actionable intents as the task that the user intended the digital assistant to perform. In some embodiments, the domain that has the most "triggered" nodes is selected. In some embodiments, the domain having the highest confidence value (e.g., based on the relative importance of its various triggered nodes) is selected. In some embodiments, the domain is selected based on a combination of the number and the importance of the triggered nodes. In some embodiments, additional factors are considered

in selecting the node as well, such as whether the digital assistant has previously correctly interpreted a similar request from a user.

**[0070]** In some embodiments, the digital assistant also stores names of specific entities in the vocabulary index 344, so that when one of these names is detected in the user request, the natural language processor 332 will be able to recognize that the name refers to a specific instance of a property or sub-property in the ontology. In some embodiments, the names of specific entities are names of businesses, restaurants, people, movies, and the like. In some embodiments, the digital assistant searches and identifies specific entity names from other data sources, such as the user's address book, a movies database, a musicians database, and/or a restaurant database. In some embodiments, when the natural language processor 332 identifies that a word in the token sequence is a name of a specific entity (such as a name in the user's address book), that word is given additional significance in selecting the actionable intent within the ontology for the user request.

**[0071]** For example, when the words "Mr. Santo" are recognized from a user request, and the last name "Santo" is found in the vocabulary index 344 as one of the contacts in the user's address book, then it is likely that the user request corresponds to a "send a message" or "initiate a phone call" domain. For another example, when the words "ABC Cafe" are found in the user request, and the term "ABC Cafe" is found in the vocabulary index 344 as the name of a particular restaurant in the user's city, then it is likely that the user request corresponds to a "restaurant reservation" or a "find restaurants" domain.

**[0072]** User data 348 includes user-specific information, such as user-specific vocabulary, user preferences, user address or home location, user's default and secondary languages, user's address book, and other short-term or long-term information for each user. In some embodiments, the natural language processor 332 uses the user-specific information to supplement the information contained in a user input to further define the user intent. For example, for a user request "invite my friends to my birthday party," the natural language processor 332 is able to access user data 348 to determine who the "friends" are and when and where the "birthday party" would be held, rather than requiring the user to provide such information explicitly in his/her request.

**[0073]** Other details of searching an ontology based on a token string is described in U.S. Utility Application Serial No. 12/341,743 for "Method and Apparatus for Searching Using

An Active Ontology," filed December 22, 2008, the entire disclosure of which is incorporated herein by reference.

**[0074]** In some implementations, natural language processor 332 includes categorization module 349. In some implementations, the categorization module 349 determines whether each of the one or more terms in a text string is one of an entity, an activity, or a location, as discussed in greater detail below. In some implementations, the categorization module 349 classifies each term of the one or more terms as one of an entity, an activity, or a location.

**[0075]** In some embodiments, once the natural language processor 332 identifies an actionable intent (or domain) based on the user request, the natural language processor 332 generates a structured query to represent the identified actionable intent. In some embodiments, the structured query includes parameters for one or more nodes within the domain for the actionable intent, and at least some of the parameters are populated with the specific information and requirements specified in the user request. For example, the user may say "Make me a dinner reservation at a sushi place at 7." In this case, the natural language processor 332 may be able to correctly identify the actionable intent to be "restaurant reservation" based on the user input. According to the ontology, a structured query for a "restaurant reservation" domain optionally includes parameters such as {Cuisine}, {Time}, {Date}, {Party Size}, and the like. In some embodiments, based on the information contained in the user's utterance, the natural language processor 332 generates a partial structured query for the restaurant reservation domain, where the partial structured query includes the parameters {Cuisine= "Sushi"} and {Time="7pm"}. In this example, the user's utterance, however, contains insufficient information to complete the structured query associated with the domain. Therefore, other necessary parameters such as {Party Size} and {Date} are not specified in the structured query based on the information currently available. In some embodiments, the natural language processor 332 populates some parameters of the structured query with received context information. For example, in some embodiments, if the user requested a sushi restaurant "near me," the natural language processor 332 populates a {location} parameter in the structured query with GPS coordinates from the user device 104.

**[0076]** In some embodiments, the natural language processor 332 passes the structured query (including any completed parameters) to the task flow processing module 336 ("task flow processor"). The task flow processor 336 is configured to receive the structured query from the natural language processor 332, complete the structured query, if necessary, and

perform the actions required to "complete" the user's ultimate request. In some embodiments, the various procedures necessary to complete these tasks are provided in task flow models 354. In some embodiments, the task flow models include procedures for obtaining additional information from the user, and task flows for performing actions associated with the actionable intent.

**[0077]** As described above, in order to complete a structured query, the task flow processor 336 sometimes needs to initiate additional dialogue with the user in order to obtain additional information, and/or disambiguate potentially ambiguous utterances. When such interactions are necessary, the task flow processor 336 invokes the dialogue processing module 334 ("dialogue processor 334") to engage in a dialogue with the user. In some embodiments, the dialogue processor 334 determines how (and/or when) to ask the user for the additional information, and receives and processes the user responses. The questions are provided to and answers are received from the users through the I/O processing module 328. In some embodiments, the dialogue processor 334 presents dialogue output to the user via audio and/or visual output, and receives input from the user via spoken or physical (e.g., clicking) responses. Continuing with the example above, when the task flow processor 336 invokes the dialogue flow processor 334 to determine the "party size" and "date" information for the structured query associated with the domain "restaurant reservation," the dialogue flow processor 335 generates questions such as "For how many people?" and "On which day?" to pass to the user. Once answers are received from the user, the dialogue flow processor 334 can then populate the structured query with the missing information, or pass the information to the task flow processor 336 to complete the missing information from the structured query.

**[0078]** In some cases, the task flow processor 336 receives a structured query that has one or more ambiguous properties, e.g., when a structured query for the "send a message" domain indicates that the intended recipient is "Bob," and the user happens to have multiple contacts named "Bob." In such cases, the task flow processor 336, optionally, requests that the dialogue processor 334 disambiguate this property of the structured query. In turn, the dialogue processor 334, optionally, asks the user "Which Bob?", and displays (or reads) a list of contacts named "Bob" from which the user may choose.

**[0079]** In some implementations, disambiguation module 350 disambiguates pronouns. In such implementations, disambiguation module 350 identifies one of the one or more terms as a pronoun and determines a noun to which the pronoun refers. In some implementations, disambiguation module 350 determines a noun to which the pronoun refers as a name of an

entity, an activity, or a location by using user data 348 or vocabulary index 344 associated with a user of the electronic device.

**[0080]** In some implementations, disambiguation module 350 accesses information obtained from one or more sensors (e.g., proximity sensor 214, light sensor 212, GPS receiver 213, temperature sensor 215, and motion sensor 210) of a handheld electronic device (e.g., user device 104) for determining a meaning of one or more of the terms. In some implementations, disambiguation module 350 identifies two terms each associated with one of an entity, an activity, or a location. For example, a first of the two terms refers to a person, and a second of the two terms refers to a location. In some implementations, disambiguation module 350 identifies three terms each associated with one of an entity, an activity, or a location.

**[0081]** Once the task flow processor 336 has completed the structured query for an actionable intent, the task flow processor 336 proceeds to perform the ultimate task associated with the actionable intent. Accordingly, the task flow processor 336 executes the steps and instructions in the task flow model according to the specific parameters contained in the structured query. For example, the task flow model for the actionable intent of "restaurant reservation", optionally, includes steps and instructions for contacting a restaurant and actually requesting a reservation for a particular party size at a particular time. For example, using a structured query such as: {restaurant reservation, restaurant = ABC Cafe, date = 3/12/2012, time = 7pm, party size = 5}, the task flow processor 336, optionally, performs the steps of: (1) logging onto a server of the ABC Cafe or a restaurant reservation system such as OPENTABLE®, (2) entering the date, time, and party size information in a form on the website, (3) submitting the form, and (4) making a calendar entry for the reservation in the user's calendar.

**[0082]** In some embodiments, the task flow processor 336 employs the assistance of a service processing module 338 ("service processor") to complete a task requested in the user input or to provide an informational answer requested in the user input. For example, the service processor 338 can act on behalf of the task flow processor 336 to make a phone call, set a calendar entry, invoke a map search, invoke or interact with other user applications installed on the user device, and invoke or interact with third party services (e.g. a restaurant reservation portal, a social networking website, a banking portal, etc.). In some embodiments, the protocols and application programming interfaces (API) required by each service can be specified by a respective service model among the services models 356. The

service processor 338 accesses the appropriate service model for a service and generates requests for the service in accordance with the protocols and APIs required by the service according to the service model.

[0083] For example, if a restaurant has enabled an online reservation service, the restaurant can submit a service mode 1 specifying the necessary parameters for making a reservation and the APIs for communicating the values of the necessary parameter to the online reservation service. When requested by the task flow processor 336, the service processor 338 can establish a network connection with the online reservation service using the web address stored in the service model, and send the necessary parameters of the reservation (e.g., time, date, party size) to the online reservation interface in a format according to the API of the online reservation service.

[0084] In some embodiments, the natural language processor 332, dialogue processor 334, and task flow processor 336 are used collectively and iteratively to infer and define the user's intent, obtain information to further clarify and refine the user intent, and finally generate a response (e.g., an output to the user, or the completion of a task) to fulfill the user's intent.

[0085] In some embodiments, after all of the tasks needed to fulfill the user's request have been performed, the digital assistant 326 formulates a confirmation response, and sends the response back to the user through the I/O processing module 328. If the user request seeks an informational answer, the confirmation response presents the requested information to the user. In some embodiments, the digital assistant also requests the user to indicate whether the user is satisfied with the response produced by the digital assistant 326.

[0086] More details on the digital assistant can be found in the U.S. Utility Application No. 12/987,982, entitled "Intelligent Automated Assistant," filed January 18, 2010, and also in U.S. Utility Application No. 61/493,201, entitled "Generating and Processing Data Items That Represent Tasks to Perform," filed June 3, 2011, the entire disclosures of which are incorporated herein by reference.

[0087] In many instances, a digital assistant is able to infer a user's intent based on a speech input in the form of a request from a user and fulfill the user's request either by providing information sought by the user request or by performing a task according to the user request. Sometimes, however, the digital assistant fails to provide a satisfactory response to the user request for information or action. The reasons for the failures can be many, such as imperfect speech recognition, unrecognized terms and concepts in the user

request, incorrect or incomplete information and inadequate capability in the digital assistant's services, and so on. Regardless of the reason for the digital assistant's failure to provide a satisfactory response to the user request, it is desirable to implement a suitable training feature for the digital assistant to prevent future failures to satisfy the user request.

**[0088]** As shown in Figure 3B, in some embodiments, the digital assistant 326 also includes a training module 340 configured to alter or adjust one or more parameters of speech-to-text processing (e.g., language models, acoustic models, vocabulary used in speech recognition), natural language processing (e.g., grammar, ontology, etc.) or task-flow processing when training module 340 detects a failure (or an impasse in a dialogue between the digital assistant and the user) to provide a satisfactory response to a user request. In some embodiments, the training module utilizes clarification inputs from the user subsequent to the failure to produce a satisfactory response to the user request. In some embodiments, the training module elicits a clarifying response from the user regarding the failed user request.

**[0089]** In some embodiments, the digital assistant 326 also includes a crowd sourcing module 342 configured to collect feedback information from one or more other DA-clients 102 (e.g., one or more external CS information sources) related to the failed user request and uses the feedback information to supplement the digital assistant's abilities in speech processing (including speech recognition), natural language processing (including intent inference), and/or task flow processing (including task execution). For example, the crowd sourced feedback information helps to recognize speech in particular accents or dialects, expand the vocabulary associated with domains, and/or identify additional domain entities. In some embodiments, the crowd sourced feedback information is used (either by the digital assistant itself, or by a provider of the digital assistant) to create additional domains and task flows to further expand the capabilities of the digital assistant.

**[0090]** In some embodiments, the crowd sourcing module 342 establishes and maintains a crowd sourced (CS) knowledge base 358. The CS knowledge base 358 stores crowd sourced information that addresses informational or task requests that the digital assistant provides to a user. In some embodiments, the contents of CS knowledge base 358 are organized by records of previous user requests to which the digital assistant had initially failed to successfully respond, but subsequently fulfilled using crowd-sourced information. CS knowledge base 358 provides references and information to the digital assistant to provide satisfactory responses to the same or similar user requests received in the future. In some embodiments, CS knowledge base 358 is organized to facilitate searching by the natural

language processor. For example, the information and answers in CS knowledge base 358 are indexed by nodes in ontology 360, so that the infrastructure of the natural language processor can be leveraged to quickly find past questions and answers in one or more relevant domains.

**[0091]** In some embodiments, the training module 340, the crowd sourcing module 342 and the CS knowledge base 358 are part of the digital assistant system 326, and can access various other components through various internal application interfaces. In some embodiments, one or more of the training module 340, the crowd sourcing module 342, and the CS knowledge base 358, or one or more sub-components thereof are optionally provided separately from the digital assistant system 326, and the digital assistant system 326 accesses each of these modules and sub-components thereof through one or more external application programming interfaces (APIs). In some embodiments, various aspects of the digital assistant system 326, such as the speech-to-text, natural language processing or task-flow processing are modified based on the information stored in the CS knowledge base 358 to improve future performance of the digital assistant system 326.

**[0092]** In some embodiments, the digital assistant maintains a user log 370 for a user of user device 104 based on user requests and interactions between the digital assistant and the user. The user log optionally stores information such as the user requests received, the context information surrounding the user requests, the responses provided to the user, and feedback provided by the user (e.g., clarification inputs or rejections), the parameters, models, third-party services, and procedures used by the digital assistant to generate and provide the responses, etc. In some embodiments, the user log is searchable by various search parameters, such as time, location, user, demographic, response type, feedback type, task type, duration, failure type, etc.

**[0093]** Although Figure 3B does not explicitly show the communication interfaces between all components of the digital assistant 326, it is to be understood that the components shown are capable of communicating with other components of the digital assistant 326 either directly or through one or more other interfaces, such as application programming interfaces (APIs), database query interfaces, and/or other interfaces, protocols, and/or communication channels.

**[0094]** Figures 4A-4C illustrate a flow diagram of a method 400 for training a digital assistant in accordance with some embodiments. In some embodiments, method 400 is

performed at an electronic device including one or more processors and memory (e.g., a respective DA-client 102 or DA-server 106). In some embodiments, method 400 is governed by a set of instructions stored in memory (e.g., a non-transitory computer readable storage media) that are executed by the one or more processors of the electronic device.

[0095] In some embodiments, during a dialogue between a digital assistant and a user (402), the electronic device receives (404) at least one speech input from the user. During a dialogue between a digital assistant and a user, Figure 3B, for example, shows I/O processing module 328 configured to receive at least one speech input from the user. The speech input, for example, includes a user request for action such as "Call Mom" or a user request for information such as "Give me directions to Big Basin State Park."

[0096] In some embodiments, during the dialogue, the electronic device infers (406) an initial intent based on the at least one speech input. Figure 3B, for example, shows speech-to-text processing module 330 configured to generate a respective token sequence associated with the user request (e.g., "Call Mom"). In turn, natural language processor 332 is configured to obtain the respective token sequence from speech-to-text processing module 330 and associate the respective token sequence with an actionable intent (or a task). This process of associating the respective token sequence with an actionable intent is sometimes referred to as intent inference. The natural language processor 332 associates the respective token sequence with an actionable intent (e.g., "initiate a phone call" actionable intent) which is associated with a domain within ontology 360 (e.g., "phone call" domain). Subsequently, the natural language processor 332 generates a structured query to represent the identified actionable intent. For example, a structured query for "initiate a phone call" optionally includes the following parameters: {name}, {phone number}, and so on. In some embodiments, the natural language processor 332 populates the {name} parameter with "Mom" from the user request and the {phone number} parameter with Mom's cell phone number stored in the user's address book in user data 348.

[0097] In some embodiments, during the dialogue, the electronic device provides (408) an initial response to fulfill the initial intent that has been inferred. Task flow processing module 336 is configured to obtain the structured query from the natural language processor 332 and perform the ultimate task associated with the actionable intent. Accordingly, the task flow processor 336 executes the steps and instructions in a respective task flow model according to the specific parameters contained in the structured query. For example, the task flow model for the actionable intent of "initiate a phone call" includes steps and instructions

for initiating a phone call. For example, using a structured query such as: {initiate a phone call, name = Mom, phone number = Mom's cell phone number}, the task flow processor 336 performs the steps of: (1) asking the user whether they wish to proceed with the phone call to Mom at her cell phone number; and (2) invoking a phone call application to call Mom on her cell phone number.

**[0098]** In some embodiments, during the dialogue, the electronic device receives (410) a follow-up speech input from the user rejecting the initial response. Figure 3B, for example, shows I/O processing module 328 configured to receive a follow-up speech input from the user rejecting the initial response. For example, the user responds to the digital assistant's initial response "Do you want me to call Mom on her cell phone?" by saying "No. Stop" (e.g., the follow-up speech input).

**[0099]** The electronic device detects (412) an impasse during the dialogue between the digital assistant and the user, where the dialogue includes at least one speech input from the user. Figure 3B, for example, shows training module 340 configured to detect an impasse during a dialogue between a digital assistant and a user (e.g., including at least one user speech input). In reference to the above example, training module 340 detects an impasse in the dialogue when the digital assistant receives a follow-up speech input from the user rejecting the initial response (e.g., "No. Stop").

**[0100]** In some embodiments, the impasse comprises one of a set consisting of (414): a single user rejection of an initial response to the at least one speech input, where the initial response is distinct from the satisfactory response; two or more user rejections of two or more corresponding distinct responses to the at least one speech input; and a user command ending the dialogue. In some implementations, the training module 340 is configured to detect an impasse after a single user rejection of an initial response as discussed in the above example. In some implementations, the training module 340 is configured to detect an impasse after two or more user rejections of two or more corresponding distinct responses to the at least one speech input. For example, the digital assistant receives a follow-up speech input from the user (e.g., "No. Stop") rejecting the initial response from the digital assistant (e.g., "Do you want me to call Mom on her cell phone?"), and, after receiving the follow-up speech input rejecting the initial response, the digital assistant asks the user "Do you want me to call Mom at work?", where the work phone number is distinct from the cell phone number. In this example, the digital assistant receives a second follow-up speech input from the user rejecting the second response (e.g., "No. I don't want to call Mom at work either."). After

receiving the second follow-up speech input from the user rejecting the second response the digital assistant detects an impasse. In some implementations, the training module 340 is configured to detect an impasse when the user executes a command exiting the dialogue with the digital assistant. For example, the user presses an end command to stop a phone call from connecting to Mom or the user locks the electronic device or the user presses a home command to exit the dialogue. In the aforementioned examples, the user interrupts completion of the task related to the user request.

**[0101]** In some embodiments, after detecting the impasse, the electronic device obtains (416) user authorization to establish the learning session. Figure 3B, for example, shows training module 340 configured to obtain authorization from the user to establish a learning session. For example, training module 340 invokes dialogue processing module 334 to query the user whether they wish to establish a learning session to cure the impasse and produce a satisfactory response to the user request.

**[0102]** In response to detecting the impasse, the electronic device establishes (418) a learning session associated with the at least one speech input from the user. Figure 3B, for example, shows training module 340 configured to establish a learning session after detecting the impasse so as to produce a satisfactory response to the user request and to prevent a future impasse associated with the user request. In some embodiments, during the learning session, training module 340 is configured to temporarily adjust one or more parameters of speech-to-text processing (including speech recognition), natural language processing (including intent inference) and/or task-flow processing (including task execution) so as to produce a satisfactory response to the user request. In some embodiments, during the learning session, training module 340 is configured to temporarily alter one or more thresholds of speech-to-text processing and/or natural language processing so as to produce more and varied responses to the user request than during normal operations.

**[0103]** In some embodiments, during the learning session and prior to receiving one or more subsequent clarification inputs from the user, the electronic device provides (420) two or more alternative responses to the least one speech input from the user. Figure 3B, for example, shows training module 340 configured to provide two or more alternative responses to the user request. For example, after the user rejected the digital assistant's initial response (e.g., "Do you want me to call Mom at her cell phone?") to the user request (e.g., "Call Mom"), the digital assistant provides the user with two alternative responses to the user request (e.g., "Do you want me to call Mom at her work phone or on Skype?"). In some

implementations, the two or more alternative responses are distinction from the initial response.

**[0104]** In some embodiments, the electronic device (424) reduces a respective intent inference or speech recognition threshold so as to generate the two or more alternative responses to the at least one speech input from the user. In some embodiments, training module 340 is configured to provide two or more alternative responses to the user request which the natural language processor 332 did not initially select due to corresponding low confidence values. After the user rejects an initial response associated with a highest confidence value, training module 340 lowers the threshold confidence level required by the natural language processor 332 in order to provide two or more alternative responses to the user. For example, the alternative responses indicated in the aforementioned example have confidence values of 65% and 40%, respectively, in comparison to an 85% confidence value for the initial response.

**[0105]** In some embodiments, during the learning session and prior to receiving one or more subsequent clarification inputs from the user, the electronic device rephrases (426) at least a portion of the at least one speech input from the user so as to elicit one or more subsequent clarification inputs from the user. For example, a user speech input is a user request for digital assistant to perform an action such as "Ring my Wife." However, the natural language processor 332 cannot infer an actionable intent from the word "ring" and an impasse occurs. In response to detecting the impasse, training module 340 is configured to rephrase the user request to elicit a clarification input from the user such as "Did you mean to call your wife?" Subsequently, in this example, the user acknowledges that this was the user's original intent by responding "Yes, I did." In turn, training module 340 alters a relationship between the nodes within ontology 360 to relate the word "ring" to the "call" property node. For example, under the "initiate a call" actionable node training module 340 adds "ring" as a synonym or a sub-property node for the "call" property node.

**[0106]** In some embodiments, the initial intent is a best guess and during the learning session, the electronic device: infers (428) a second intent based on the at least one speech input, prior to receiving the one or more subsequent clarification inputs from the user, where the second intent is a second best guess and the second intent is distinct from the initial intent; and provides a second response to fulfill the second intent that has been inferred. For example, in response to a user request (e.g., "Call Mom"), the digital assistant provides the best guess of the initial intent to the user by asking the user "Do you want me to call Mom on

her cell phone?" In some implementations, the natural language processor 332 initially selects the best guess, where the ontology indicates a highest confidence value associated with the best guess. In response to the initial response, the user replies "No. Stop" and training module 340 detects an impasse in the dialogue. Subsequently, during the learning session, training module 340 provides the second best guess to the user such as "Do you want me to call Mom at work?" In some implementations, the natural language processor 332 subsequently selects the second best guess, where the ontology indicates a second highest confidence value associated with the second best guess.

[0107] During the learning session, the electronic devices receives (430) one or more subsequent clarification inputs from the user. In some embodiments, a clarification input includes a user speech input repeating the user request or initial speech input from the user, a user speech input repeating the user request or initial speech input from the user with a changed emphasis, a user speech input spelling out a word misinterpreted or misidentified by the digital assistant, a user speech input accepting an alternative response of two or more alternative responses provided by the digital assistant, a user speech input confirming the rephrased user request provided by the digital assistant, a user speech input accepting a second best guess provided by the digital assistant, a user speech input further expounding upon or clarifying the initial user request, and so on. For example, after a user request to perform an action such as "Send a message to Thorn Yorke" and the digital assistant initially responding by saying "No such name in address book," the user repeats their initial request by saying "Send a message to Thorn Yorke." In another example, after a user request to perform an action such as "Call Phillip Giese" (pronounced giz-ay) and the digital assistant initially responds by saying "No such name in address book," the user clarifies their request by spelling out the contact's last name (e.g., "I know it's in there. Call Phillip Giz-ay. G-T E-S-E"). In another example, after a user request for information such as "Give me directions to Big Basin State Park" and the digital assistant initially responding by displaying directions to Big Basin State Park from the user's home location stored in user data 348, the user clarifies their initial request by saying "No. Give me directions to Big Basin State Park from my current location." In another example, after a user request to perform an action such as "Tell FaceSpace to set my status to 'Radiohead concert tonight!'" and the digital assistant initially responds by asking the user "Do you want me to send a message to FaceSpace saying 'Radiohead concert tonight?'" the user clarifies their initial request by stating "No, I don't want you to send a message to anyone. Set my status in my FaceSpace app to 'Radiohead

concert tonight!" In each of the aforementioned examples, the follow-up speech input from the user is a clarification input.

**[0108]** In some embodiments, the one or more subsequent clarification inputs comprise (432) one or more user text inputs. In some embodiments, a clarification input from a user includes a user text input spelling out a word misinterpreted or misidentified by the digital assistant or a user text input repeating the initial vocal user request. For example, after a user request such as "Find me the closest Italian restaurant" and the digital assistant initially responds by stating "I have found 3 stores selling Italian shoes nearby," the user types "Italian food" via a keyboard displayed on the touch-screen of device 104. In another example, after a user request to perform an action such as "Call Jimmy Hoffa" and the digital assistant initially responds by asking the user "Do you want to call Jimmy Hofford?", the user types "C-A-L-L H-O-F-F-A" via a keyboard displayed on the touch-screen of device 104.

**[0109]** In some embodiments, the one or more subsequent clarification inputs comprise (434) a definition or a pronunciation of the at least one speech input. For example, a clarification input from a user includes a user speech input defining a word misinterpreted or misidentified by the digital assistant or a user speech input emphasizing the correct pronunciation of a word misinterpreted or misidentified by the digital assistant. For example, after a user request for information such as "Find me the nearest Charcuterie store" and the digital assistant initially responds by stating "I do not know what that word means," the user provides a definition for Charcuterie. For example, after a user request for information such as "Give me the projected temperature for tomorrow in Columbia" and the digital assistant initially responds by displaying tomorrow's weather forecast for Bogota, Colombia, the user clarifies their request by emphasizing the pronunciation of Columbia (e.g., "No. I want tomorrow's temperature for Co-LUM-bee-a South Carolina").

**[0110]** During the learning session, the electronic device adjusts (436) at least one of intent inference and task execution associated with the at least one speech input so as to produce a satisfactory response to the speech input based at least in part on the one or more subsequent clarification inputs. In some embodiments, training module 340 is configured to adjust at least one of speech-to-text processing (e.g., the language models, acoustic models or vocabulary used in speech recognition), natural language processing (e.g., the grammar or ontology used in intent inference) and task-flow processing (e.g., the task flow models used in task execution) so as to produce a satisfactory response based at least in part on the one or more clarification inputs provided by the user. For example, training module 340 is

configured to adjust a task flow model associated with initiating a phone call (e.g., altering a task flow to call Mom at work instead of at her cell phone number) or a task flow associated with sending a message (e.g., setting a status in FaceSpace instead of sending a message to FaceSpace when the user requests to set a status in FaceSpace) after a clarification input by the user. In another example, training module 340 is configured to adjust a parameter of the natural language processor 332 by adding a synonym to a property node in ontology 360 (e.g., add "ring" as a synonym to "call"), adjusting the relationship between nodes in ontology 360, or altering a domain associated with a user request in ontology 360 (e.g., associating the speech input "Find me a table" with a "furniture" domain instead of a "restaurant reservation" domain). In another example, training module 340 is configured to adjust a language or speech model associated with speech recognition (e.g., associating "Columbia" with "co-LUM-bee-a" instead of "ko-LOM-bee-a").

**[0111]** In some embodiments, the satisfactory response to the at least one speech input does not elicit (438) any further rejection by the user. Figure 3B, for example, shows training module 340 configured to determine that a response to a user request is a satisfactory response when the response does not elicit a subsequent rejection from the user. For example, after the digital assistant displays a list of nearby Indian restaurants in response to a user request for information such as "Find me a nearby Indian restaurant," the user selects one of the restaurants in the list to read a review of the selected restaurant, the user selects a restaurant in the list to call and/or the user goes to a restaurant in the list (either immediately or sometime in the future). In this example, training module 340 determines that the displayed list of Indian restaurants is a satisfactory response to the user request. In another example, after the digital assistant makes a reservation at a Ramen restaurant nearby the user's home at 8:00 PM Saturday for 2 in response to a user request for action such "Make me a reservation for 2 at a Ramen place for this coming Saturday night," the user does not object to the reservation and/or the user goes to the Ramen restaurant on Saturday night. In this example, training module 340 determines that the restaurant reservation is a satisfactory response to the user request.

**[0112]** During the learning session, the electronic device associates (440) the satisfactory response with the at least one speech input for processing future occurrences of the at least one speech input. Figure 3B, for example, shows the training module 340 configured to associate a satisfactory response with the user request (e.g., including at least one speech input) by permanently (or semi-permanently) adjusting at least one of speech-to-text

processing, natural language processing and task-flow processing so that future occurrences of the user request or speech input produce the satisfactory response instead of an impasse. In this way, a respective initial response to the at least one speech input is replaced with the satisfactory response. In some embodiments, a user is enabled to modify or revert the adjustments made by the digital assistant by initiating a dialogue with the digital assistant or going to a website associated with the user's digital assistant.

**[0113]** In some embodiments, associating the satisfactory response with the at least one speech input comprises replacing (442) a respective initial response to the at least one speech input with the satisfactory response for a set of users in a community of users. In some embodiments, DA-server 106 stores global (or regional or cultural, etc.) speech-to-text processing models (e.g., the language models, acoustic models and vocabulary used in speech recognition), natural language processing data (e.g., the grammar and one or more ontologies used in intent inference) or task-flow processing models (e.g., the task flows used in task execution) shared by a plurality of DA-clients 102. In some embodiments, in response to producing a satisfactory response to a user request, training module 340 provides the satisfactory response to the speech input to DA-server 106. In turn, DA-server 106 adjusts at least one of the globally (or regionally) shared speech-to-text processing models, natural language processing data or task-flow processing models so as to replace an initial response to the speech input with the satisfactory response to the speech input produced by training module 340. In some embodiments, DA-server 106 stores a plurality of ontologies. For example, DA-server 106 stores a city-wide ontology shared by a set of users in the community of users that have a home city in common. In this example, the user's satisfactory response is used to alter the city-wide ontology associated with the user.

**[0114]** In some embodiments, the electronic device: receives (444) respective clarification inputs from a set of users in a community of users associated with the at least one speech input; and adjusts (446) at least one of speech recognition, intent inference and task execution associated with the at least one speech input to produce the satisfactory response to the at least speech input for the user during the learning session based at least in part on the respective clarification inputs from the set of users in the community of users. Figure 3B, for example, shows crowd sourcing module 342 configured to obtain clarification inputs and/or satisfactory responses (sometimes called crowd sourced feedback information) associated with a respective user request (e.g. including at least one speech input):from other DA-clients 102. In some implementations, all DA-clients 102 within operating environment 100 are

assembled in a community of users. In some embodiments, crowd sourcing module 342 obtains crowd sourced feedback information from a set of users in the community of users. For example, crowd sourcing module 342 obtains crowd sourced feedback information from one or more sets of users in the community of users that share a commonality with the user such as the user's home city (e.g., Sunnyvale), the user's home metropolitan area (e.g., San Francisco Bay Area), the user's home state (e.g., California), the user's home region (e.g., West Coast), the user's nation/country (e.g., U.S.A.), the user's job (e.g., engineer), the user's alma mater (e.g., Clemson University), the user's ethnicity (e.g., Uyghur), and other such commonalities. Figure 3B, for example, shows CS Knowledge base 358 configured to store crowd sourced feedback information obtained by crowd sourcing module 342. Training module 340 is configured to utilize the crowd sourced feedback information associated with a respective user request to adjust at least one of speech-to-text processing (including speech recognition), national language processing (including intent inference) and task-flow processing (including task execution) so as to produce a satisfactory response to the respective user request.

**[0115]** It should be understood that the particular order in which the operations in Figures 4A-4C have been described is merely exemplary and is not intended to indicate that the described order is the only order in which the operations could be performed. One of ordinary skill in the art would recognize various ways to reorder the operations described herein.

**[0116]** Figure 5 illustrates a block diagram of training module 340 in accordance with some embodiments. Training module 340 includes a pattern identification module 510, a hypothesis generation module 512, a monitoring module 514, a parameter alteration module 516, an analytics module 518, and a generalization module 520. Pattern identification module 514 is configured to identify a pattern of success or failure as to the completion of a respective task based on the feedback information collected from a plurality of events each of which are associated with a previous completion of the respective task. In some embodiments, the feedback information collected from the plurality of events is stored in user log 370. In some embodiments, pattern identification module 510 is further configured to identify a pattern of success or failure as to the completion of a respective task based on the feedback information in user log 370 and crowd sourced feedback information from at least a set of DA-clients 102 in a community of users associated with the completion of the respective task.

[0117] Hypothesis generation module 512 is configured to generate a hypothesis regarding a parameter associated with at least one of speech-to-text processing (including speech recognition), natural language processing (including intent inference) and task-flow processing (including task execution) as a cause for the pattern of success or failure identified by pattern identification module 510. Monitoring module 514 is configured to monitor speech-to-text processing and natural language processing to identify subsequent requests by the user to complete the respective task or a task similar to the respective task. Parameter alteration module 516 is configured to temporarily adjust at least one of speech-to-text processing, natural language processing or task-flow processing to test the hypothesis generated by hypothesis generation module 512. Analytics module 518 is configured to adopt or reject the hypothesis generated by hypothesis generation module 512 based on feedback information collected from subsequent completions of the respective task stored in user log 370. Optionally, generalization module 520, if included in training module 340, is configured to provide the hypothesis to DA-server 106 so as to test the hypothesis on one or more sets of DA-clients 102 in the community of users. In some embodiments, each of the one or more sets of DA-clients 102 in the community of users share a commonality (e.g., location, region, culture, accent, dialect, etc.) with the user.

[0118] Although Figure 5 does not explicitly show the communication interfaces between all components of the digital assistant 326 in Figure 3B, it is to be understood that the components shown in Figure 5 are capable of communicating with other components of the digital assistant 326 in Figure 3B either directly or through one or more other interfaces, such as application programming interfaces (APIs), database query interfaces, and/or other interfaces, protocols, and/or communication channels.

[0119] Figures 6A-6C illustrate a flow diagram of a method 600 for training a digital assistant in accordance with some embodiments. In some embodiments, method 600 is performed at an electronic device including one or more processors and memory (e.g., a respective DA-client 102 or DA-server 106). In some embodiments, method 600 is performed at the digital assistant 326 illustrated in Figure 5. In some embodiments, method 600 is governed by a set of instructions stored in memory (e.g., a non-transitory computer readable storage media) that are executed by the one or more processors of the electronic device.

[0120] The electronic device obtains (602) feedback information associated with one or more previous completions of a task. After completing a user request (also sometimes called

a task) for action (e.g., calling a contact or sending a message to a contact or setting a reminder to pick up dry cleaning) or information (e.g., displaying directions to a synagogue or displaying a list of haberdasheries), the digital assistant 326 is configured to collect feedback information associated with the completion of the task and store the feedback information in user log 370. Digital assistant 326 collects both explicit and implicit feedback information regarding the completion of a task. Explicit feedback information, for example, includes a user speech input rejecting a set of displayed directions, a user command exiting the dialogue with the digital assistant (e.g., pressing the home button, turning off the device, or locking the device), a user speech input affirming a suggestion by the digital assistant, and so on. Implicit feedback information, for example, includes a user physically going to a restaurant suggested by the digital assistant, a user following the directions suggested by the digital assistant, a user not physically going to a restaurant suggested by the digital assistant, and so on. Digital assistant 326 collects feedback information in the form of a speech input from the user, a text input from the user, user actions, context information (e.g., information pertaining to the hardware or software state of the user device), and so on.

**[0121]** In some embodiments, the feedback information comprises (604) a user rejection of a response associated with a completion of a task. For example, after a user request for information such as "Find me Italian food" and the digital assistant responds by displaying nearby Italian restaurants, the user rejects the displayed results by saying "No. I want Italian gelato" Or, in another example, after a user request to perform an action such as "Call Christina" and the digital assistant responds by calling Christina at her work phone number, the user responds by saying "No. Stop. Call Christina at her cell phone number."

**[0122]** In some embodiments, the feedback information comprises (606) a user location subsequent to the completion of a task. For example, after a user request for information such as "Find me a concert in San Francisco on Saturday night" and the digital assistant responds by displaying a list of concerts in San Francisco for this coming Saturday night (including a Less than Jake concert at The Independent), the digital assistant detects the user's GPS coordinates via user device 104 at The Independent in San Francisco on Saturday night.

**[0123]** In some embodiments, the feedback information comprises (608) a length of time spent at a location nearby a completed task. For example, after a user request for information such as "Find me a nearby museum" and the digital assistant responds by displaying a list of museums close to the user's home location including the Tech Museum of Innovation in San

Jose and the Rosicrucian Egyptian Museum in San Jose, the digital assistant detects the user's GPS coordinates via user device 104 at the Tech Museum for 10 minutes and at the Rosicrucian Egyptian Museum for 120 minutes. In some embodiments, the amount of time that a user spends at a location associated with a completed task affects the indicator value corresponding to the completed event (indicator values are discussed at length below).

**[0124]** In some embodiments, the feedback information comprises (610) a user action subsequent to a completion of a task. For example, after a user request for information such as "Give me the Caltrain schedule for today" and the digital assistant responds by displaying today's Northbound timetable for Caltrain, the digital assistant detects the user on a train headed north between Sunnyvale and Palo Alto via the acoustic environment, GPS coordinates, network signals, etc. detected by user device 104.

**[0125]** The electronic device identifies (612) a pattern of success or failure associated with an aspect of speech recognition, intent inference or task execution previously used to complete the task. Figure 5, for example, shows pattern identification module 510 configured to identify a pattern of success or failure associated with an aspect of speech recognition, intent inference and task execution previously used to complete the task based on feedback information stored in user log 370. The feedback information included in user log 370 comprises a plurality of indicators of success or failure associated with previous completions of a respective task. In some embodiments, each of the indicators of success or failure are associated with an indicator value between 1 and 0, where an indicator value close to 1 denotes a strong indicator of success and an indicator value close to 0 denotes a strong indicator of failure. For example, user log 370 stores a pool of feedback information related to a "find restaurants" user request (or task), where each entry of feedback information relates to an event in which the "find restaurants" actionable intent node was triggered. In this example, pattern identification module 510 identifies a pattern of failure related to a respective user request such as "Find me a Greek restaurant" due to the fact that user log 370 includes 3 feedback information entries with an indicator of success (e.g., an indicator value greater than 0.5) and 5 feedback entries with an indicator of failure (e.g., an indicator value less than or equal to 0.5) related to the user request. In relation to indicator values, for example, if a user goes to a cobbler on a list of cobblers displayed in response to the user's request (e.g., "Find me a cobbler"), the indicator value associated with the feedback information entry will equal 0.95 if the user goes to the cobbler within 5 minutes of the

displayed list of cobblers or the indicator value will equal 0.55 if the user goes to the cobbler two weeks after the displayed list of cobblers.

**[0126]** In some embodiments, identifying a pattern of success or failure comprises (614) detecting a learning opportunity, where detecting a learning opportunity comprises identifying one or more previous failures of completion of the task in a user log. Figure 5, for example, shows pattern identification module 510 configured to detect a learning opportunity (i.e., an opportunity to improve at least one of speech recognition, intent inference or task execution) by identifying a pattern of failure in user log 370 from one or more previous indicators of failure associated with a respective task. In some embodiments, pattern identification module 510 identifies a learning opportunity when the normalized sum of the indicators of failure outweigh the normalized sum of the indicators of success for a respective task or a ratio of indicators of failure to indicators of success for a respective task is equal to or greater than 1.

**[0127]** In some embodiments, the electronic device provides (616) to the user the one or more previous failures of completion of the task to elicit a subsequent clarification from the user. For example, after pattern identification module 510 identifies a learning opportunity related to a respective task, the digital assistant provides one or more of the previous failures to complete the respective task to the user to elicit a subsequent clarification so as to improve at least one of speech recognition, intent inference or task execution related to the respective task. For example, pattern identification module 510 identifies a learning opportunity related to a user request (also sometimes called a task) for information such as "Find me a bar" because the pool of events in user log 370 related to this task indicates a pattern of failure. In this example, digital assistant seizes upon the learning opportunity by asking the user "When you say 'Find me a bar' do you want me to find dive bars, dance clubs or lounge bars?" in order to elicit a subsequent clarification response from the user related to the respective task so as to improve at least one of speech recognition, intent inference or task execution related to the respective task.

**[0128]** In some embodiments, digital assistant is configured to actively seize upon a learning opportunity when a user opts into a training mode (e.g., a user says "Let's clear up some misunderstandings between us") or a user self-identifies him or herself as a problematic speaker or digital assistant detects that a user is idle (e.g., if a user is simply bantering with the digital assistant, the digital assistant asks the user whether he or she wishes to enter a training mode).

**[0129]** The electronic device generates (618) a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution as a cause for the pattern of success or failure. Figure 5, for example, shows hypothesis generation module 512 configured to generate a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution as a cause for the pattern of success or failure identified by pattern identification module 510. For example, pattern identification module 510 identifies a pattern of failure regarding completion of a task related to calling a contact within the user's address book (e.g., "Call Phillip Giese") and, in turn, hypothesis generation module 512 generates a hypothesis regarding a parameter used in speech recognition (e.g., the speech or language model for the contact's last name should be modified) as a cause for the pattern of failure. In another example, pattern identification module 510 identifies a pattern of failure regarding completion of a task related to a verb not found in ontology 360 (e.g., "Ring my wife.") and in turn, hypothesis generation module 512 generates a hypothesis regarding a parameter used in intent inference (e.g., "ring" should be added as a synonym or sub-property node for "call" within ontology 360) as a cause for the pattern of failure. In some embodiments, a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution includes a hypothesis that adjusting a language or acoustic model for speech-to-text processing, adding a word to the vocabulary for speech-to-text processing, altering a relationship between nodes or domains in ontology 360 for natural language processing, adding property nodes to ontology 360 for natural language processing, or altering a task flow model for task flow processing, would improve upon the pattern of success or failure identified by pattern identification module 510.

**[0130]** The electronic device identifies (620) one or more subsequent requests for completion of the task. Figure 5, for example, shows monitoring module 514 configured to identify one or more subsequent requests by the user for completion of a task. For example, monitoring module 514 is configured to monitor each token sequence produced by speech-to-text processing module 330 and/or each structured query produced by the natural language processor 332 to identify a subsequent request for completion of the task.

**[0131]** In some embodiments, after identifying a respective subsequent request for completion of the task, the electronic device obtains (622) user authorization to test the hypothesis. In some embodiments, digital assistant is configured to obtain user authorization before testing the hypothesis as to the cause of the pattern of success or failure for a task. For example, training module 340 invokes dialogue processing module 334 to query the user

whether they wish to test a hypothesis related to a task (e.g., a user request to perform a task) so as to produce a satisfactory response to the user request.

**[0132]** The electronic device tests (624) the hypothesis by altering the parameter used in the at least one of speech recognition, intent inference and task execution for subsequent completions of the task. Figure 5, for example, shows parameter alteration module 516 configured to test the hypothesis generated by hypothesis generation module 512 by temporarily altering a parameter used in the at least one of speech recognition, intent inference and task execution for a subsequent request for completion of a task detected by monitoring module 516.

**[0133]** In some embodiments, testing the hypothesis occurs (626) when a hypothesis confidence value associated with the hypothesis exceeds a predetermined confidence threshold. In some embodiments, hypothesis generation module 512 is configured to assign a confidence value between 0 and 1 to all hypotheses it generates based on the feedback information associated with the one or more previous completions of the task associated with the hypothesis. A confidence value close to 0 denotes a low likelihood of the hypothesis being the cause of the pattern of success or failure, and a confidence value close to 1 denotes a high likelihood of the hypothesis being the cause of the pattern of success or failure. In practice, however, a hypothesis confidence value never equals 0 or 1. Furthermore, in some embodiments, parameter alteration module 516 is configured to test the hypothesis by altering a parameter used in the at least one of speech recognition, intent inference and task execution for a subsequent completion of the task when the confidence value assigned to the hypothesis exceeds a predetermined confidence threshold such as 0.5, 0.67, or 0.75.

**[0134]** The electronic device adopts or rejects (628) the hypothesis based on feedback information collected from the subsequent completions of the task. For example, analytics module 518 (Figure 5) is configured to adopt or reject the hypothesis based on feedback information collected from the one or more subsequent completions of the task. In some embodiments, when analytics module 518 adopts the hypothesis, parameter alteration module 516 is configured to permanently or semi-permanently alter a parameter used in the at least one of speech recognition, intent inference and task execution for subsequent completions of the task.

**[0135]** In some embodiments, adopting or rejecting the hypothesis comprises (630) determining whether a task completion metric associated with the task has improved. In

some embodiments, a task completion metric is a ratio of indicators of failure to indicators of success based on feedback information related to completions of a respective task (or user request). In some embodiments, analytics module 618 is configured to adopt a hypothesis if the task completion metric improves. In other words, analytics module 618 is configured to adopt the hypothesis if after implementing the hypothesis, the ratio of indicators of failures to indicators of success for the task is less than the ratio prior to implementing the hypothesis (e.g., the task completion metric improves as it approaches 0).

**[0136]** In some embodiments, testing the hypothesis comprises (632) performing a test on a first user and after adopting the hypothesis based on feedback information collected from the subsequent completions of the task by the first user, the electronic device: tests (634) the hypothesis on a set of one or more users in a community of users, where the first user is among the set of the one or more users in the community of users; and adopts or rejects (636) the hypothesis based on feedback information collected from the subsequent completions of the task by the set of the one or more users in the community of users, where adopting the hypothesis includes implementing the hypothesis for at least a portion of the set of the one or more users in the community of users. For example, parameter alteration module 516 is configured to test a hypothesis generated by hypothesis generation module 512 on a first user by temporarily altering a parameter used in at least one of speech recognition, intent inference and task execution for a subsequent completions of the task detected. In accordance with a determination that analytics module 518 adopts the hypothesis, generalization module 520 is configured to further generalize the hypothesis to a set of one or more users in a community of users, where the first user is among the set of users, and to provide the hypothesis to DA-server 106.

**[0137]** In some embodiments, the first user and the set of one or more users in the community of users share a commonality. For example, both the user and the set of users live in Sunnyvale, CA or both the user and the set of users are nuclear engineers. In turn, DA-server 106 tests the hypothesis on the set of users by temporarily altering a parameter used in at the least one of speech recognition, intent inference and task execution for subsequent completions of the task by the set of users. DA-server 106 adopts or rejects the hypothesis based on feedback information collected from the subsequent completions of the task by the set of users in the community of users. DA-server 106 adopts the hypothesis by permanently or semi-permanently altering the parameter used in at the least one of speech recognition, intent inference and task execution for at least a portion of the set of one or more

users in the community of users. In some embodiments, DA-server 106 further attempts to generalize the hypothesis by testing the hypothesis against a super set of users among which the set of users belongs or against the entire community of users.

**[0138]** In some embodiments, after rejecting (638) the hypothesis, the electronic device: generates (640) a modified hypothesis regarding the parameter used in at least one of speech recognition, intent deduction, and task execution as the cause for the pattern of success or failure; and tests (642) the modified hypothesis by altering the parameter in at least one of speech recognition, intent deduction, and task execution for subsequent completions of the task. Figure 5, for example, shows hypothesis generation module 512 configured to generate a modified hypothesis in view of an original hypothesis after the analytics module 520 rejects the original hypothesis. For example, after analytics module 520 rejects an original hypothesis with a 0.8 confidence value related to a task because the completion metric did not improve, hypothesis generation module is configured to generate a modified hypothesis related to the task with a lower confidence value. In this example, parameter alteration module 516 is configured to test the modified hypothesis by temporarily altering the parameter in at least one of speech recognition, intent deduction, and task execution for subsequent completions of the task. For example, a modified hypothesis includes a more drastic (or less drastic) alteration to a speech model in speech recognition or a relationship between nodes for intent inference (natural language processing) in comparison to the first (or original) hypothesis that was rejected by analytics module 518.

**[0139]** It should be understood that the particular order in which the operations in Figures 6A-6C have been described is merely exemplary and is not intended to indicate that the described order is the only order in which the operations could be performed. One of ordinary skill in the art would recognize various ways to reorder the operations described herein.

**[0140]** In accordance with some embodiments, Figure 7 shows a functional block diagram of an electronic device 700 configured in accordance with the principles of the various described embodiments. The functional blocks of the device are, optionally, implemented by hardware, software, or a combination of hardware and software to carry out the principles of the various described embodiments. It is understood by persons of skill in the art that the functional blocks described in Figure 7 are, optionally, combined or separated into sub-blocks to implement the principles of the various described embodiments. Therefore, the description

herein optionally supports any possible combination or separation or further definition of the functional blocks described herein.

**[0141]** As shown in Figure 7, an electronic device 700 includes a sound receiving unit 702 configured to receive sound input, a speaker unit 704 configured to output sound, optionally, a touch-screen display unit 706 configured to display information and to receive touch contacts; and a processing unit 708 coupled to the sound receiving unit 702, the speaker unit 704 and, optionally, the touch-screen display unit 706. In some embodiments, the processing unit 708 includes a detecting unit 710, an establishing unit 712, an adjusting unit 714, and an associating unit 716.

**[0142]** The processing unit 708 is configured to detect (e.g., with the detecting unit 710) an impasse during a dialogue between the digital assistant and a user, where the dialogue includes at least one speech input (e.g., received with the sound receiving unit 702) from the user. In response to detecting the impasse, the processing unit 708 is configured to establish (e.g., with the establishing unit 712) a learning session associated with the at least one speech input. During the learning session, the processing unit 708 is configured to: receive (e.g., with sound receiving unit 702 or the touch-screen display unit 706) one or more subsequent clarification inputs from the user; based at least in part on the one or more subsequent clarification inputs, adjust (e.g., with the adjusting unit 714) at least one of intent inference and task execution associated with the at least one speech input to produce a satisfactory response to the at least one speech input; and associate (e.g., with the associating unit 716) the satisfactory response with the at least one speech input for processing future occurrences of the at least one speech input.

**[0143]** In accordance with some embodiments, Figure 8 shows a functional block diagram of an electronic device 800 configured in accordance with the principles of the various described embodiments. The functional blocks of the device are, optionally, implemented by hardware, software, or a combination of hardware and software to carry out the principles of the various described embodiments. It is understood by persons of skill in the art that the functional blocks described in Figure 8 are, optionally, combined or separated into sub-blocks to implement the principles of the various described embodiments. Therefore, the description herein optionally supports any possible combination or separation or further definition of the functional blocks described herein.

**[0144]** As shown in Figure 8, an electronic device 800 includes a sound receiving unit 802 configured to receive sound input, a speaker unit 804 configured to output sound, optionally, a touch-screen display unit 806 configured to display information and to receive touch contacts; and a processing unit 808 coupled to the sound receiving unit 802, the speaker unit 804 and, optionally, the touch-screen display unit 806. In some embodiments, the processing unit 808 includes an obtaining unit 810, an identifying unit 812, a generating unit 814, a monitoring unit 816, an altering unit 818, and an analytics unit 820.

**[0145]** The processing unit 808 is configured to: obtain (e.g., with the obtaining unit 810) feedback information associated with one or more previous completions of a task; and identify (e.g., with the identifying unit 812) a pattern of success or failure associated with an aspect of speech recognition, intent inference or task execution previously used to complete the task. The processing unit 808 is further configured to: generate (e.g., with the generating unit 814) a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution as a cause for the pattern of success or failure; and identify (e.g., with the monitoring unit 816) one or more subsequent requests for completion of the task. The processing unit 808 is further configured to: test the hypothesis by altering (e.g., with the altering unit 818) the parameter used in the at least one of speech recognition, intent inference and task execution for subsequent completions of the task; and adopt or reject (e.g., with the analytics unit 820) the hypothesis based on feedback information collected from the subsequent completions of the task.

**[0146]** The foregoing description, for purpose of explanation, has been described with reference to specific implementations. However, the illustrative discussions above are not intended to be exhaustive or to limit the disclosed implementations to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The implementations were chosen and described in order to best explain the principles and practical applications of the disclosed ideas, to thereby enable others skilled in the art to best utilize them with various modifications as are suited to the particular use contemplated.

**[0147]** It will be understood that, although the terms "first," "second," etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element from another. For example, a first sound detector could be termed a second sound detector, and, similarly, a second sound detector could be termed a first sound detector, without changing the meaning of the description, so long as all occurrences of the "first sound detector" are renamed consistently

and all occurrences of the "second sound detector" are renamed consistently. The first sound detector and the second sound detector are both sound detectors, but they are not the same sound detector.

**[0148]** The terminology used herein is for the purpose of describing particular implementations only and is not intended to be limiting of the claims. As used in the description of the implementations and the appended claims, the singular forms "a", "an" and "the" are intended to include the plural forms as well, unless the context clearly indicates otherwise. It will also be understood that the term "and/or" as used herein refers to and encompasses any and all possible combinations of one or more of the associated listed items. It will be further understood that the terms "comprises" and/or "comprising," when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof.

**[0149]** As used herein, the term "if" may be construed to mean "when" or "upon" or "in response to determining" or "in accordance with a determination" or "in response to detecting," that a stated condition precedent is true, depending on the context. Similarly, the phrase "if it is determined [that a stated condition precedent is true]" or "if [a stated condition precedent is true]" or "when [a stated condition precedent is true]" may be construed to mean "upon determining" or "upon a determination that" or "in response to determining" or "in accordance with a determination" or "upon detecting" or "in response to detecting" that the stated condition precedent is true, depending on the context.

What is claimed is:

1. A computer-implemented method for training a digital assistant, performed at an electronic device including one or more processors and memory storing instructions for execution by the one or more processors, the method comprising:

detecting an impasse during a dialogue between the digital assistant and a user, wherein the dialogue includes at least one speech input from the user;

in response to detecting the impasse, establishing a learning session associated with the at least one speech input;

during the learning session:

receiving one or more subsequent clarification inputs from the user; based at least in part on the one or more subsequent clarification inputs,

adjusting at least one of intent inference and task execution associated with the at least one speech input to produce a satisfactory response to the at least one speech input; and

associating the satisfactory response with the at least one speech input for processing future occurrences of the at least one speech input.

2. The method of claim 1, wherein detecting the impasse comprises:

and during the dialogue between the digital assistant and the user:

receiving at least one speech input from the user;

inferring an initial intent based on the at least one speech input;

providing an initial response to fulfill the initial intent that has been inferred;

receiving a follow-up speech input from the user rejecting the initial response.

3. The method of claim 2, wherein the initial intent is a best guess; and during the learning session, the method further comprising:

prior to receiving the one or more subsequent clarification inputs from the user, inferring a second intent based on the at least one speech input, wherein the second intent is a second best guess and the second intent is distinct from the initial intent; and

providing a second response to fulfill the second intent that has been inferred.

4. The method of any of claims 1-3, wherein the impasse comprises one of a set consisting of:

a single user rejection of an initial response distinct from the satisfactory response;  
two or more user rejections of the initial response; and  
a user command ending the dialogue.

5. The method of any of claims 1-4, further comprising:  
after detecting the impasse, obtaining user authorization to establish the learning session.
6. The method of any of claims 1-2 and 4-5, further comprising:  
during the learning session:  
prior to receiving the one or more subsequent clarification inputs from the user, providing two or more alternative responses to the at least one speech input from the user.
7. The method of claim 5, further comprising:  
during the learning session:  
reducing a respective intent inference or speech recognition threshold so as to generate the two or more alternative responses to the at least one speech input from the user.
8. The method of any of claims 1-2 and 4-5, further comprising:  
during the learning session:  
prior to receiving the one or more subsequent clarification inputs from the user, rephrasing at least a portion of the at least one speech input from the user to elicit one or more subsequent clarification inputs from the user.
9. The method of any of claims 1-8, wherein the one or more clarification inputs comprise one or more user text inputs subsequent to the at least one speech input.
10. The method of any of claims 1-8, wherein the one or more clarification inputs comprise a definition or a pronunciation of the at least one speech input.
11. The method of any of claims 1-10, wherein the satisfactory response to the at least one speech input does not elicit any further rejection by the user.

12. The method of any of claims 1-11, wherein associating the satisfactory response with the speech input comprises replacing a respective initial response to the at least speech input with the satisfactory response for a set of users in a community of users.

13. The method of any of claims 1-12, further comprising:  
receiving respective clarification inputs from a set of users in a community of users associated with the at least one speech input; and  
based at least in part on the respective clarification inputs from the set of users in the community of users, adjusting at least one of speech recognition, intent inference and task execution associated with the at least one speech input to produce the satisfactory response to the at least speech input for the user during the learning session.

14. A computer-implemented method for training a digital assistant, performed at an electronic device including one or more processors and memory storing instructions for execution by the one or more processors, the method comprising:

during a dialogue between the digital assistant and a user:  
receiving an initial speech input from a user;  
inferring an initial intent based on the initial speech input;  
providing an initial response to fulfill the initial intent that has been inferred;  
receiving a follow-up speech input rejecting the initial response;  
upon receiving the follow-up speech rejecting the initial response, establishing a learning session associated with the initial speech input;  
during the learning session:  
adjusting at least one of intent inference and task execution associated with the initial speech input to produce a satisfactory response to the initial speech input; and  
associating the satisfactory response with the initial speech input for processing future occurrences of the initial speech input.

15. A computer-implemented method for training a digital assistant, performed at an electronic device including one or more processors and memory storing instructions for execution by the one or more processors, the method comprising:

obtaining feedback information associated with one or more previous completions of a task;

identifying a pattern of success or failure associated with an aspect of speech recognition, intent inference or task execution previously used to complete the task;  
generating a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution as a cause for the pattern of success or failure;  
identifying one or more subsequent requests for completion of the task;  
testing the hypothesis by altering the parameter used in the at least one of speech recognition, intent inference and task execution for subsequent completions of the task; and  
adopting or rejecting the hypothesis based on feedback information collected from the subsequent completions of the task.

16. The method of claim 15, wherein feedback information comprises a user rejection of a response associated with a completion of a task.

17. The method of claim 15, wherein feedback information comprises a user location subsequent to a completion of a task.

18. The method of claim 15, wherein feedback information comprises a length of time spent at a location nearby a completed task.

19. The method of claim 15, wherein feedback information comprises a user action subsequent to a completion of a task.

20. The method of any of claims 15-19, further comprising:  
after identifying a respective subsequent request for completion of the task, obtaining user authorization to test the hypothesis.

21. The method of any of claims 15-20, wherein testing the hypothesis occurs when a hypothesis confidence value associated with the hypothesis exceeds a predetermined confidence threshold.

22. The method of any of claims 15-21, wherein adopting or rejecting the hypothesis based on feedback information collected from the subsequent completions of the task comprises determining whether a task completion metric associated with the task has improved.

23. The method of any of claims 15-22, wherein testing the hypothesis comprises performing a test on a first user; and

the method further comprising:

after adopting the hypothesis based on feedback information collected from the subsequent completions of the task by the first user, testing the hypothesis on a set of one or more users in a community of users, wherein the user is among the first set of the one or more users in the community of users; and

adopting or rejecting the hypothesis based on feedback information collected from the subsequent completions of the task by the set of the one or more users in the community of users, wherein adopting the hypothesis includes implementing the hypothesis for at least a portion of the set of the one or more users in the community of users.

24. The method of any of claims 15-22, further comprising:

after rejecting the hypothesis, generating a modified hypothesis regarding the parameter used in at least one of speech recognition, intent deduction, and task execution as the cause for the pattern of success or failure; and

testing the modified hypothesis by altering the parameter in at least one of speech recognition, intent deduction, and task execution for subsequent completions of the task.

25. The method of any of claims 15-24, wherein:

identifying a pattern of success or failure comprises detecting a learning opportunity; and

detecting a learning opportunity comprises identifying one or more previous failures of completion of the task in a user log.

26. The method of any of claims 25, further comprising:

providing to the user the one or more previous failures of completion of the task to elicit a subsequent clarification from the user.

27. A portable electronic device or a computer system, comprising:

one or more processors; and

memory storing one or more programs for execution by the one or more processors, the one or more programs including instructions for performing any of the methods of claims 1-26.

28. A computer readable storage medium storing one or more programs, the one or more programs comprising instructions, which, when executed by a portable electronic device or a computer system with one or more processors, cause the device to perform any of the methods of claims 1-26.

29. A graphical user interface on a portable electronic device or a computer system with a display, a memory, and one or more processors to execute one or more programs stored in the memory, the graphical user interface comprising user interfaces displayed in accordance with any of the methods of claims 1-26.

30. A portable electronic device or a computer system, comprising:  
means for performing any of the methods of claims 1-26.

31. An information processing apparatus for use in a portable electronic device or a computer system, comprising:  
means for performing any of the methods of claims 1-26.

32. A portable electronic device or a computer system, comprising:  
a processing unit configured to perform any of the methods of claims 1-26.

33. An electronic device, comprising:  
a sound receiving unit configured to receive sound input; a speaker unit configured to output sound; and  
a processing unit coupled to the sound receiving unit and the speaker unit, the processing unit configured to:

detect an impasse during a dialogue between the digital assistant and a user, wherein the dialogue includes at least one speech input from the user;

in response to detecting the impasse, establish a learning session associated with the at least one speech input;

during the learning session:

receive one or more subsequent clarification inputs from the user;  
based at least in part on the one or more subsequent clarification

inputs, adjust at least one of intent inference and task execution associated with the at least one speech input to produce a satisfactory response to the at least one speech input; and

associate the satisfactory response with the at least one speech input for processing future occurrences of the at least one speech input.

34. The electronic device of claim 33, wherein detecting the impasse comprises: during the dialogue between the digital assistant and the user, the processing unit is further configured to:

receive at least one speech input from the user;  
infer an initial intent based on the at least one speech input;  
provide an initial response to fulfill the initial intent that has been inferred; and  
receive a follow-up speech input from the user rejecting the initial response.

35. The electronic device of claim 34, wherein the initial intent is a best guess; and during the learning session, the processing unit is further configured to:

prior to receiving the one or more subsequent clarification inputs from the user, infer a second intent based on the at least one speech input, wherein the second intent is a second best guess and the second intent is distinct from the initial intent; and  
provide a second response to fulfill the second intent that has been inferred.

36. The electronic device of any of claims 33-35, wherein the impasse comprises one of a set consisting of:

a single user rejection of an initial response distinct from the satisfactory response;  
two or more user rejections of the initial response; and  
a user command ending the dialogue.

37. The electronic device of any of claims 33-36, wherein the processing unit is further configured to:

after detecting the impasse, obtain user authorization to establish the learning session.

38. The electronic device of any of claims 33-34 and 36-37, wherein during the learning session, the processing unit is further configured to:

prior to receiving the one or more subsequent clarification inputs from the user, provide two or more alternative responses to the at least one speech input from the user.

39. The electronic device of claim 38, wherein during the learning session, the processing unit is further configured to:

reduce a respective intent inference or speech recognition threshold so as to generate the two or more alternative responses to the at least one speech input from the user.

40. The electronic device of any of claims 33-34 and 36-37, wherein during the learning session, the processing unit is further configured to:

prior to receiving the one or more subsequent clarification inputs from the user, rephrase at least a portion of the at least one speech input from the user to elicit one or more subsequent clarification inputs from the user.

41. The electronic device of any of claims 33-40, wherein the one or more clarification inputs comprise one or more user text inputs subsequent to the at least one speech input.

42. The electronic device of any of claims 33-40, wherein the one or more clarification inputs comprise a definition or a pronunciation of the at least one speech input.

43. The electronic device of any of claims 33-42, wherein the satisfactory response to the at least one speech input does not elicit any further rejection by the user.

44. The electronic device of any of claims 33-43, wherein associating the satisfactory response with the speech input comprises replacing a respective initial response to the at least speech input with the satisfactory response for a set of users in a community of users.

45. The electronic device of any of claims 33-44, wherein the processing unit is further configured to:

receive respective clarification inputs from a set of users in a community of users associated with the at least one speech input; and

based at least in part on the respective clarification inputs from the set of users in the community of users, adjust at least one of speech recognition, intent inference and task execution associated with the at least one speech input to produce the satisfactory response to the at least speech input for the user during the learning session.

46. An electronic device, comprising:

a sound receiving unit configured to receive sound input; a speaker unit configured to output sound; and

a processing unit coupled to the sound receiving unit and the speaker unit, the processing unit configured to:

during a dialogue between the digital assistant and a user:

receive an initial speech input from a user;

infer an initial intent based on the initial speech input;

provide an initial response to fulfill the initial intent that has been inferred;

receive a follow-up speech input rejecting the initial response;

upon receiving the follow-up speech rejecting the initial response, establish a learning session associated with the initial speech input; during the learning session:

adjust at least one of intent inference and task execution associated with the initial speech input to produce a satisfactory response to the initial speech input; and

associate the satisfactory response with the initial speech input for processing future occurrences of the initial speech input.

47. An electronic device, comprising:

a sound receiving unit configured to receive sound input; a speaker unit configured to output sound; and

a processing unit coupled to the sound receiving unit and the speaker unit, the processing unit configured to:

obtain feedback information associated with one or more previous completions of a task;

identify a pattern of success or failure associated with an aspect of speech recognition, intent inference or task execution previously used to complete the task;

generate a hypothesis regarding a parameter used in at least one of speech recognition, intent inference and task execution as a cause for the pattern of success or failure;

identify one or more subsequent requests for completion of the task;

test the hypothesis by altering the parameter used in the at least one of speech recognition, intent inference and task execution for subsequent completions of the task; and

adopt or reject the hypothesis based on feedback information collected from the subsequent completions of the task.

48. The electronic device of claim 47, wherein feedback information comprises a user rejection of a response associated with a completion of a task.

49. The electronic device of claim 47, wherein feedback information comprises a user location subsequent to a completion of a task.

50. The electronic device of claim 47, wherein feedback information comprises a length of time spent at a location nearby a completed task.

51. The electronic device of claim 47, wherein feedback information comprises a user action subsequent to a completion of a task.

52. The electronic device of any of claims 47-51, wherein the processing unit is further configured to:

after identifying a respective subsequent request for completion of the task, obtain user authorization to test the hypothesis.

53. The electronic device of any of claims 47-52, wherein the processing unit is configured to test the hypothesis when a hypothesis confidence value associated with the hypothesis exceeds a predetermined confidence threshold.

54. The electronic device of any of claims 47-53, wherein the processing unit is configured to adopt or reject the hypothesis when a task completion metric associated with the task has improved.

55. The electronic device of any of claims 47-54, wherein testing the hypothesis comprises performing a test on a first user; and

the processing unit is further configured to:

after adopting the hypothesis based on feedback information collected from the subsequent completions of the task by the first user, test the hypothesis on a set of one or

more users in a community of users, wherein the user is among the first set of the one or more users in the community of users; and

adopt or reject the hypothesis based on feedback information collected from the subsequent completions of the task by the set of the one or more users in the community of users, wherein adopting the hypothesis includes implementing the hypothesis for at least a portion of the set of the one or more users in the community of users.

56. The electronic device of any of claims 47-54, wherein the processing unit is further configured to:

after rejecting the hypothesis, generate a modified hypothesis regarding the parameter used in at least one of speech recognition, intent deduction, and task execution as the cause for the pattern of success or failure; and

test the modified hypothesis by altering the parameter in at least one of speech recognition, intent deduction, and task execution for subsequent completions of the task.

57. The electronic device of any of claims 47-56, wherein:  
identifying a pattern of success or failure comprises detecting a learning opportunity;  
and  
detecting a learning opportunity comprises identifying one or more previous failures of completion of the task in a user log.

58. The electronic device of any of claim 57, wherein the processing unit is further configured to:

provide to the user the one or more previous failures of completion of the task to elicit a subsequent clarification from the user.

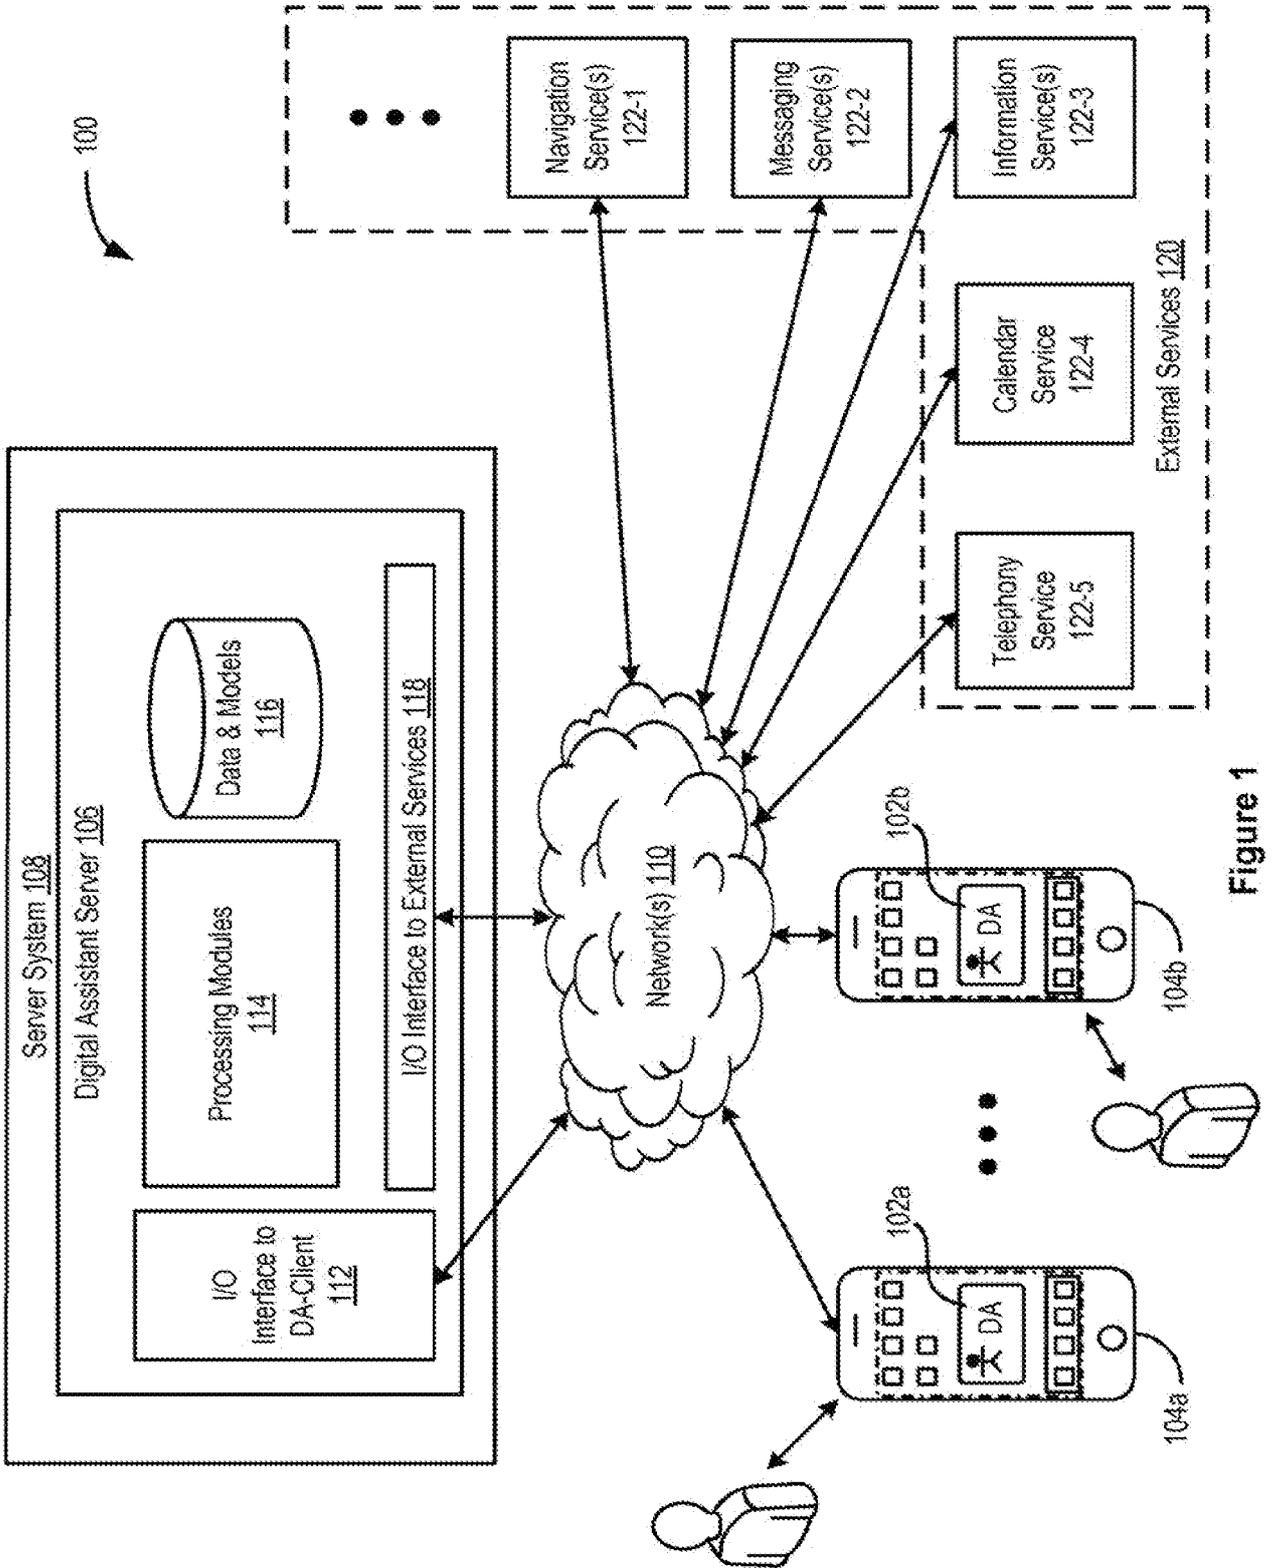


Figure 1

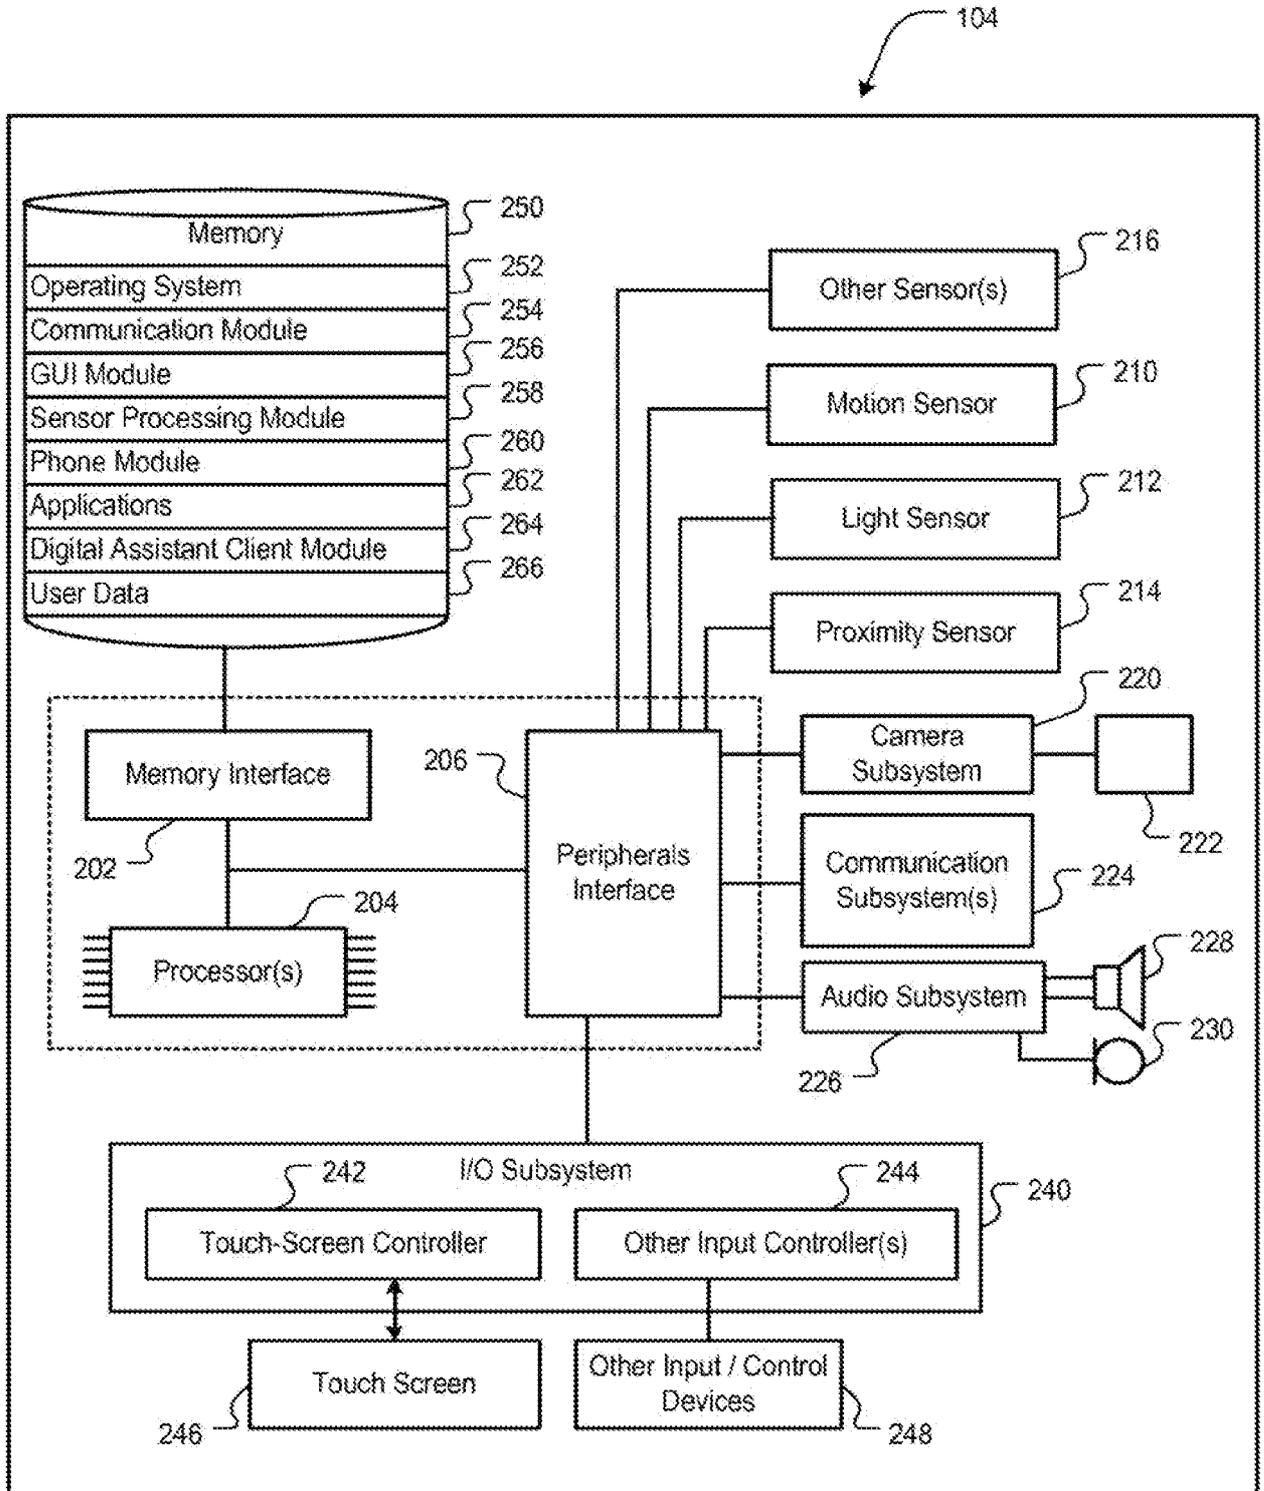


Figure 2

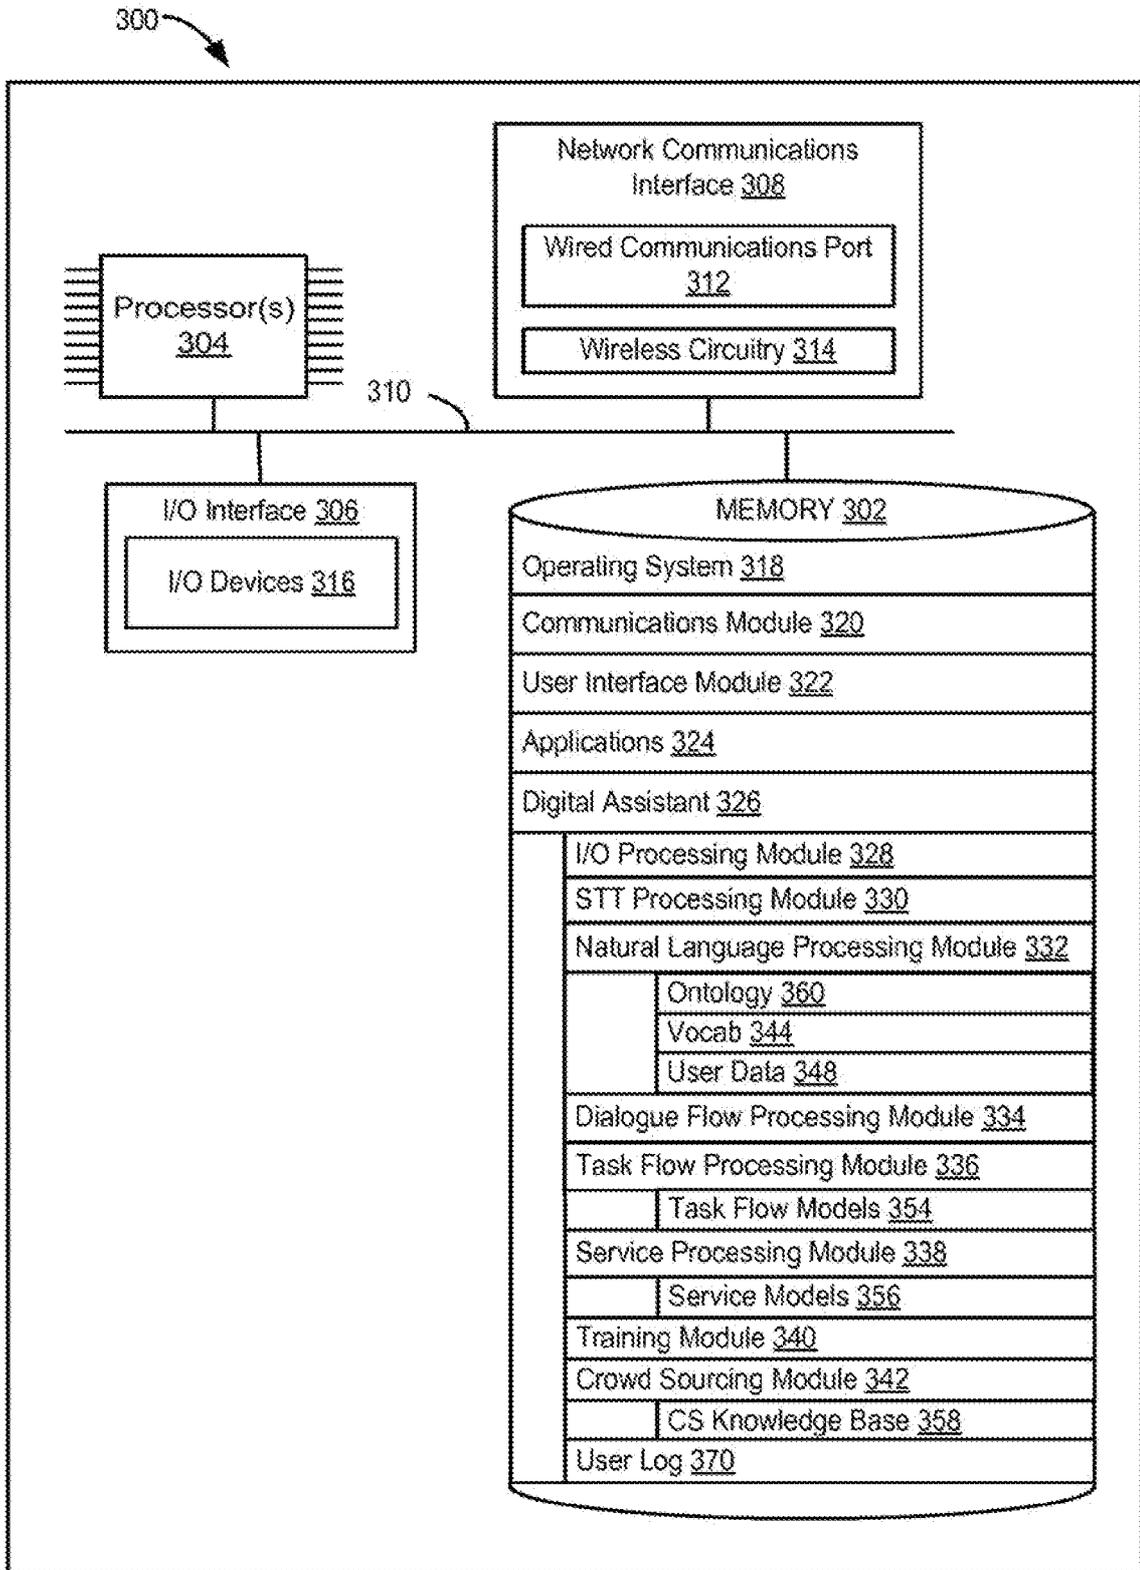


Figure 3A

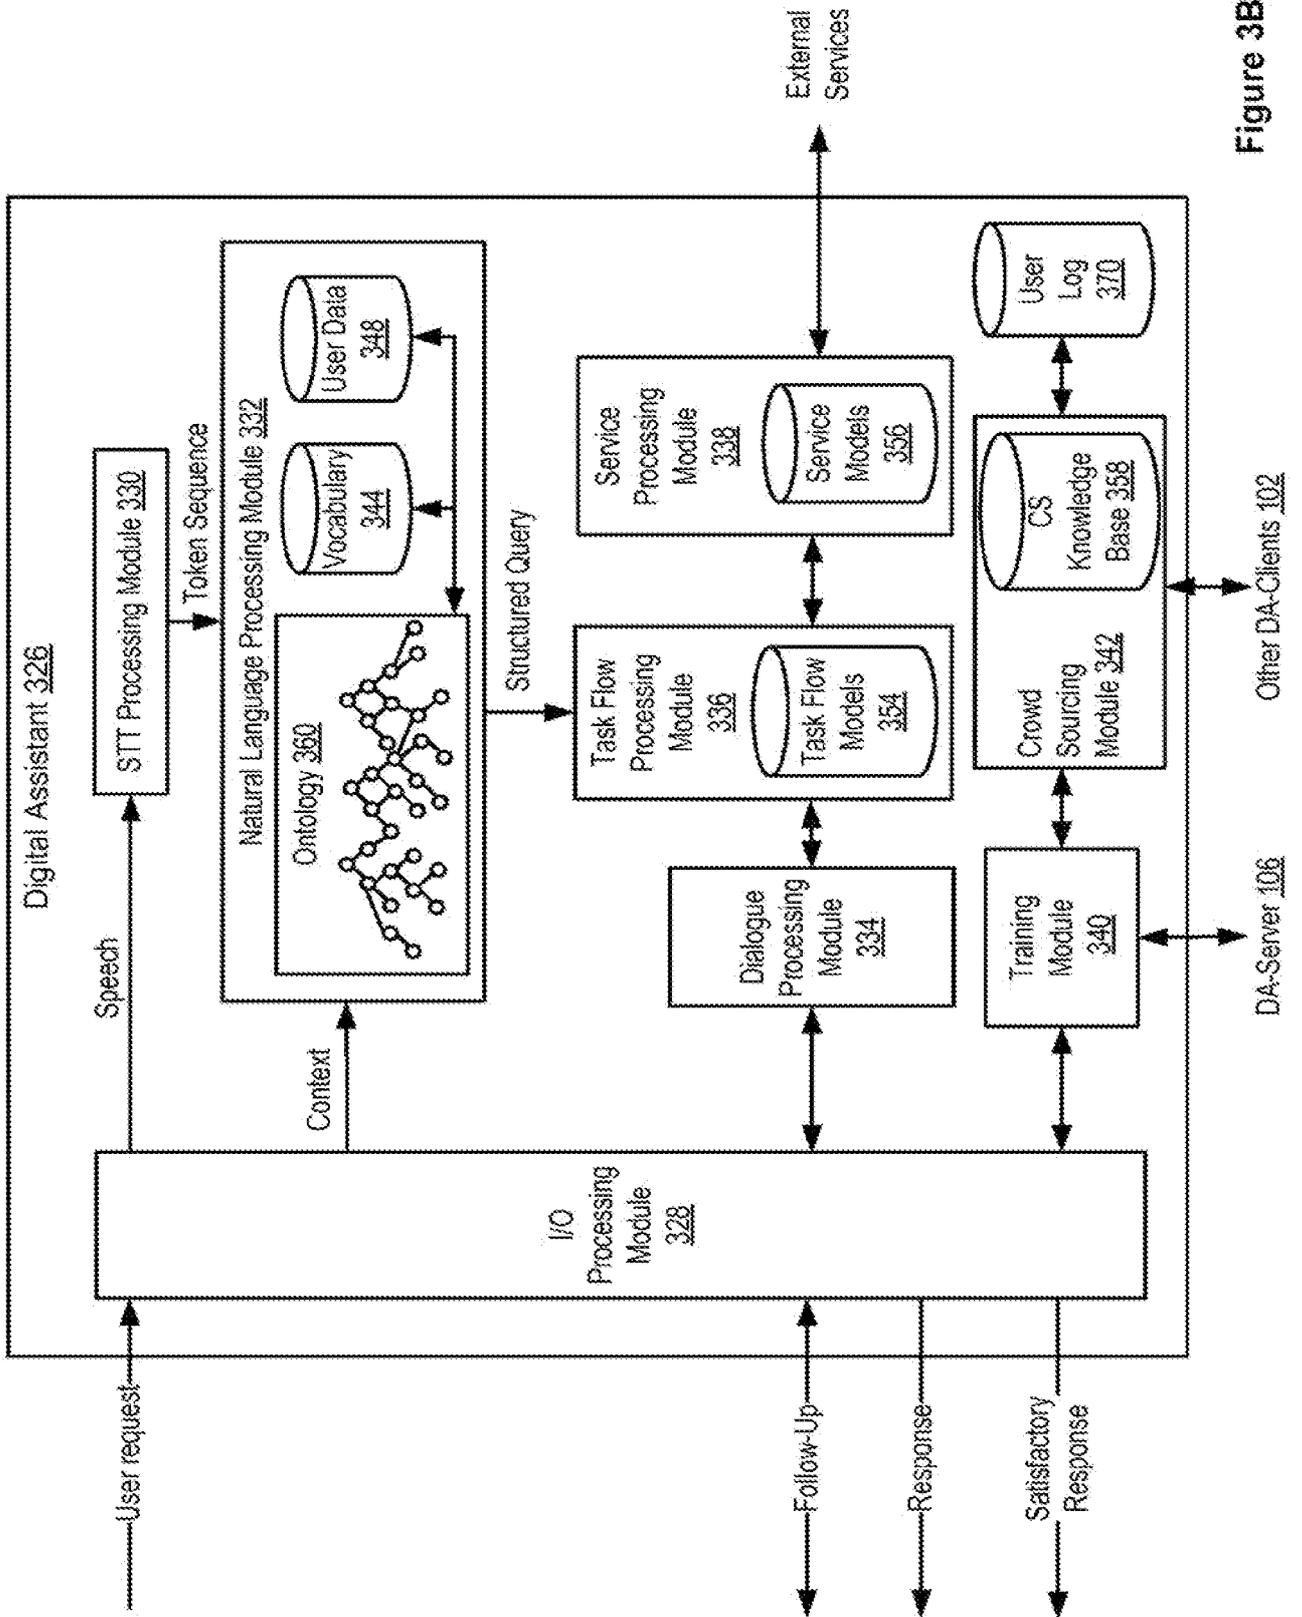


Figure 3B

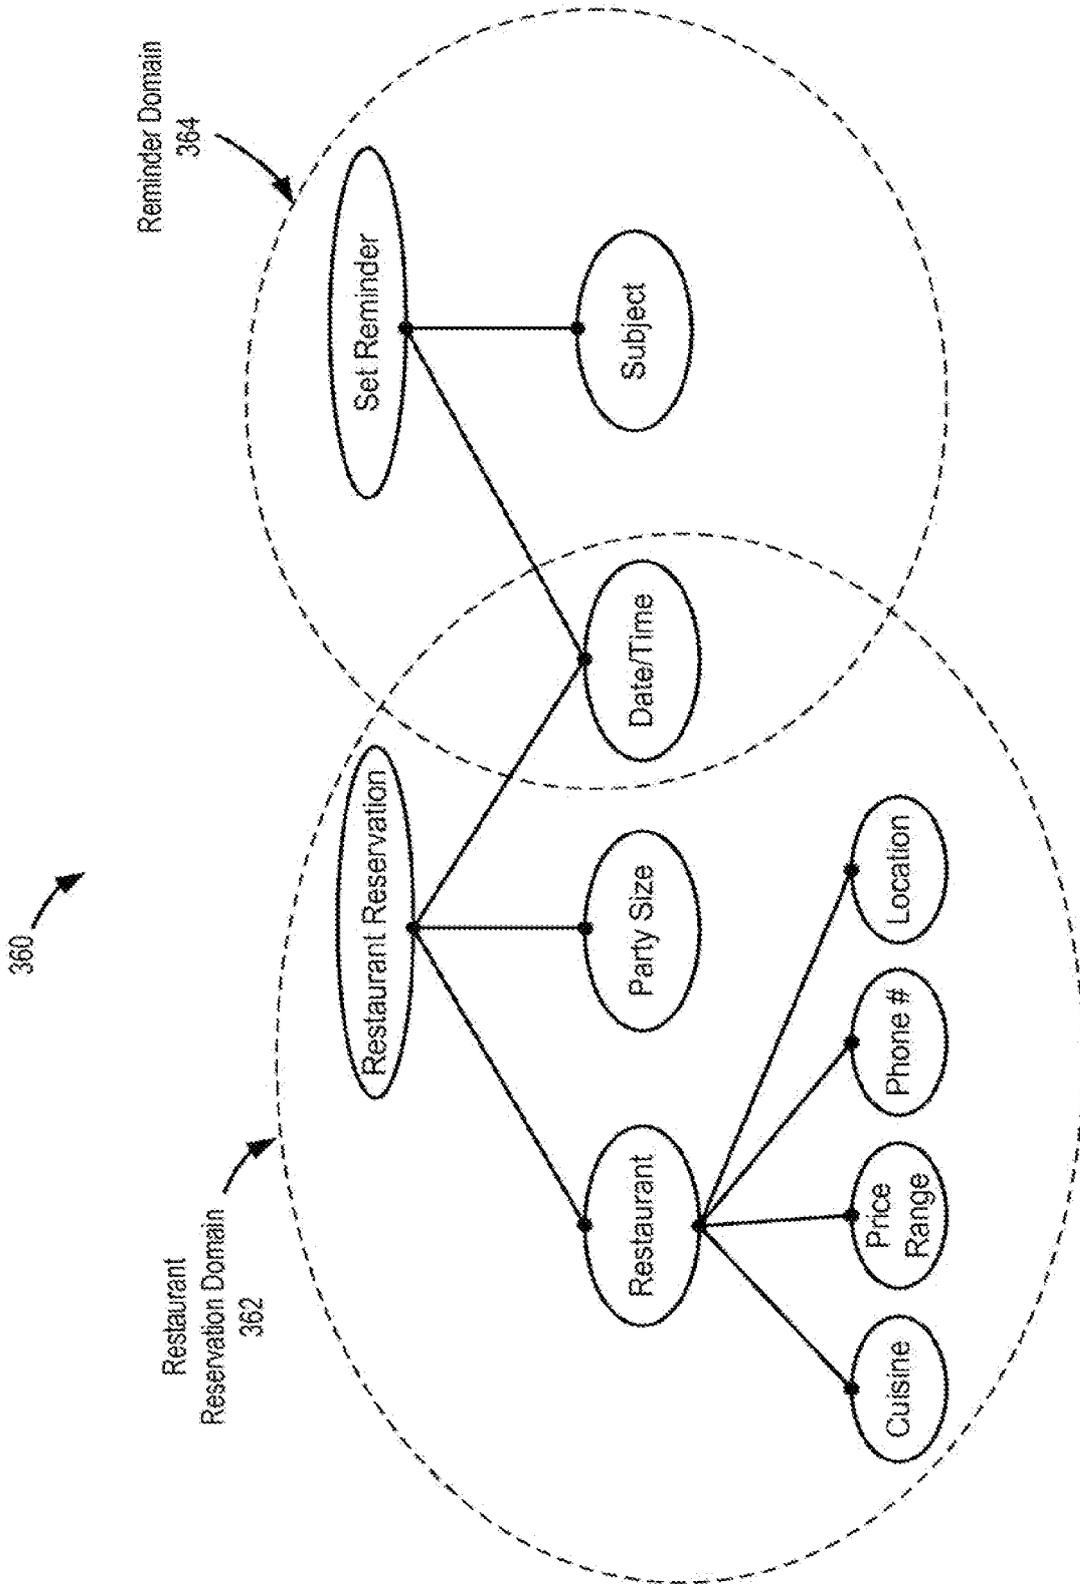


Figure 3C

400

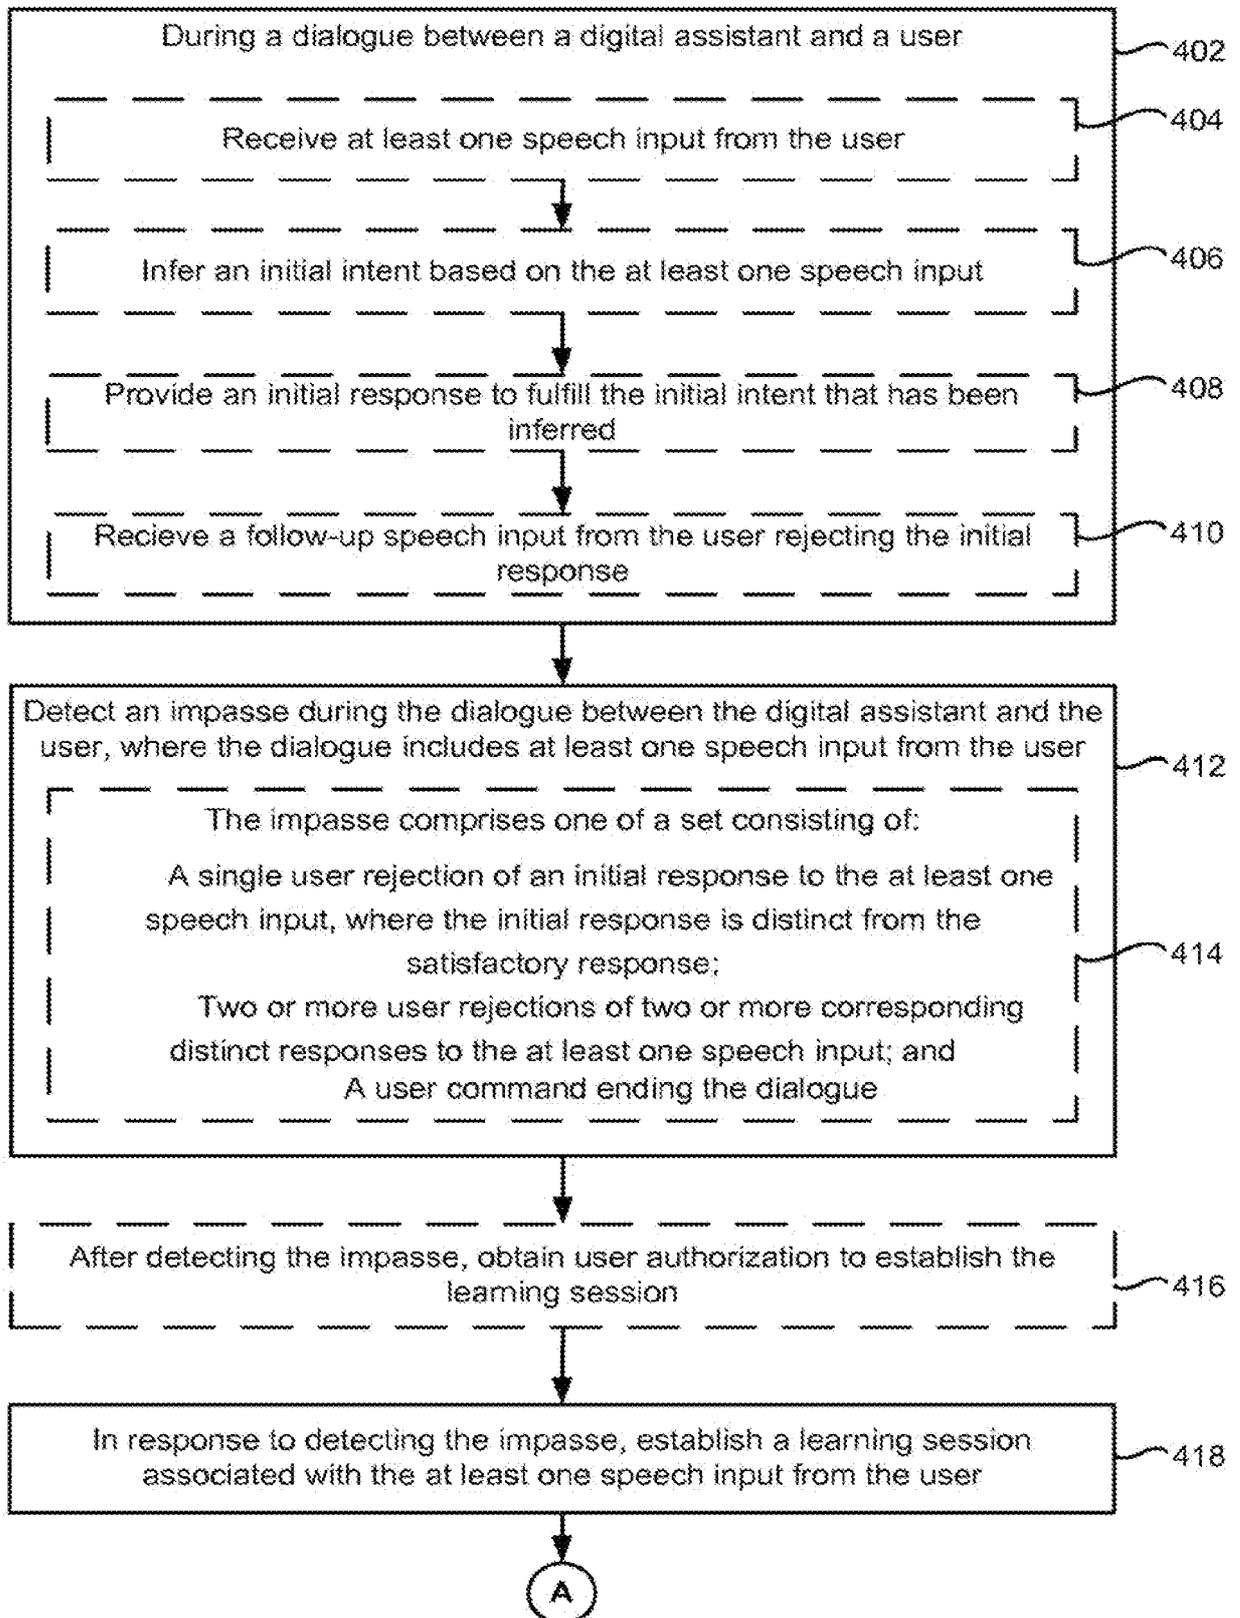


Figure 4A

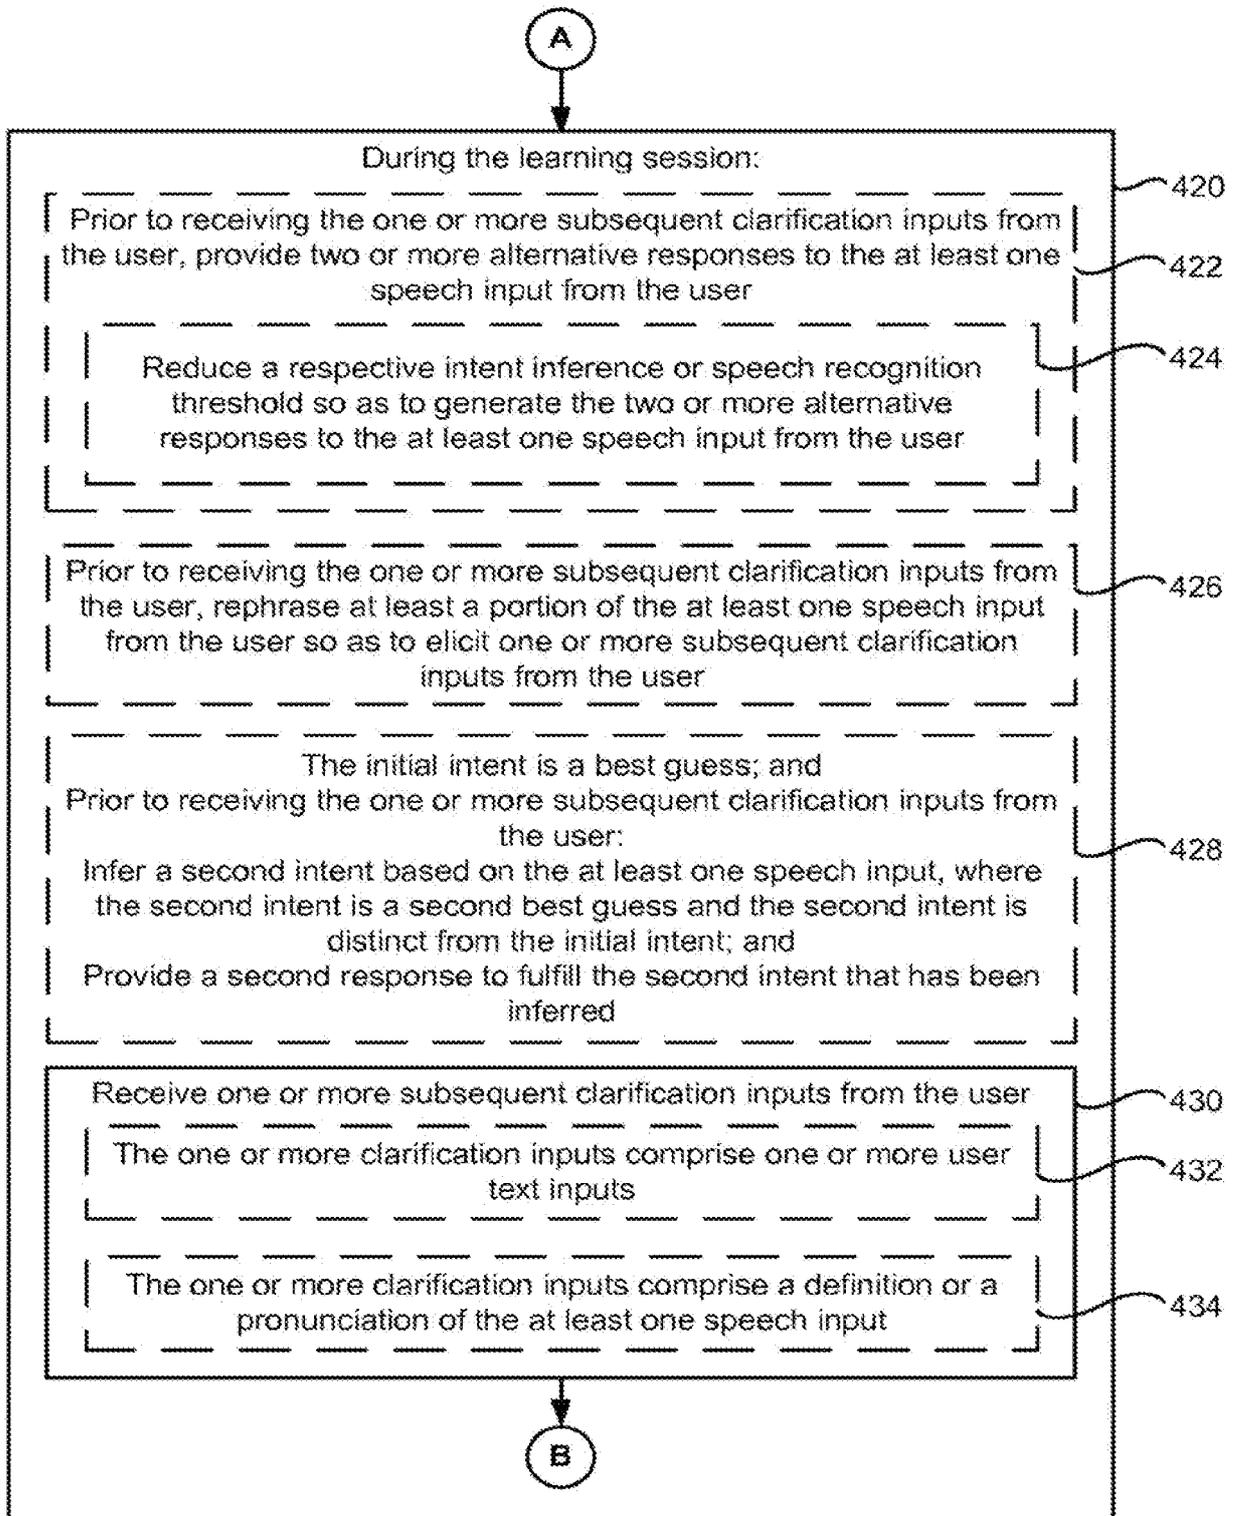


Figure 4B

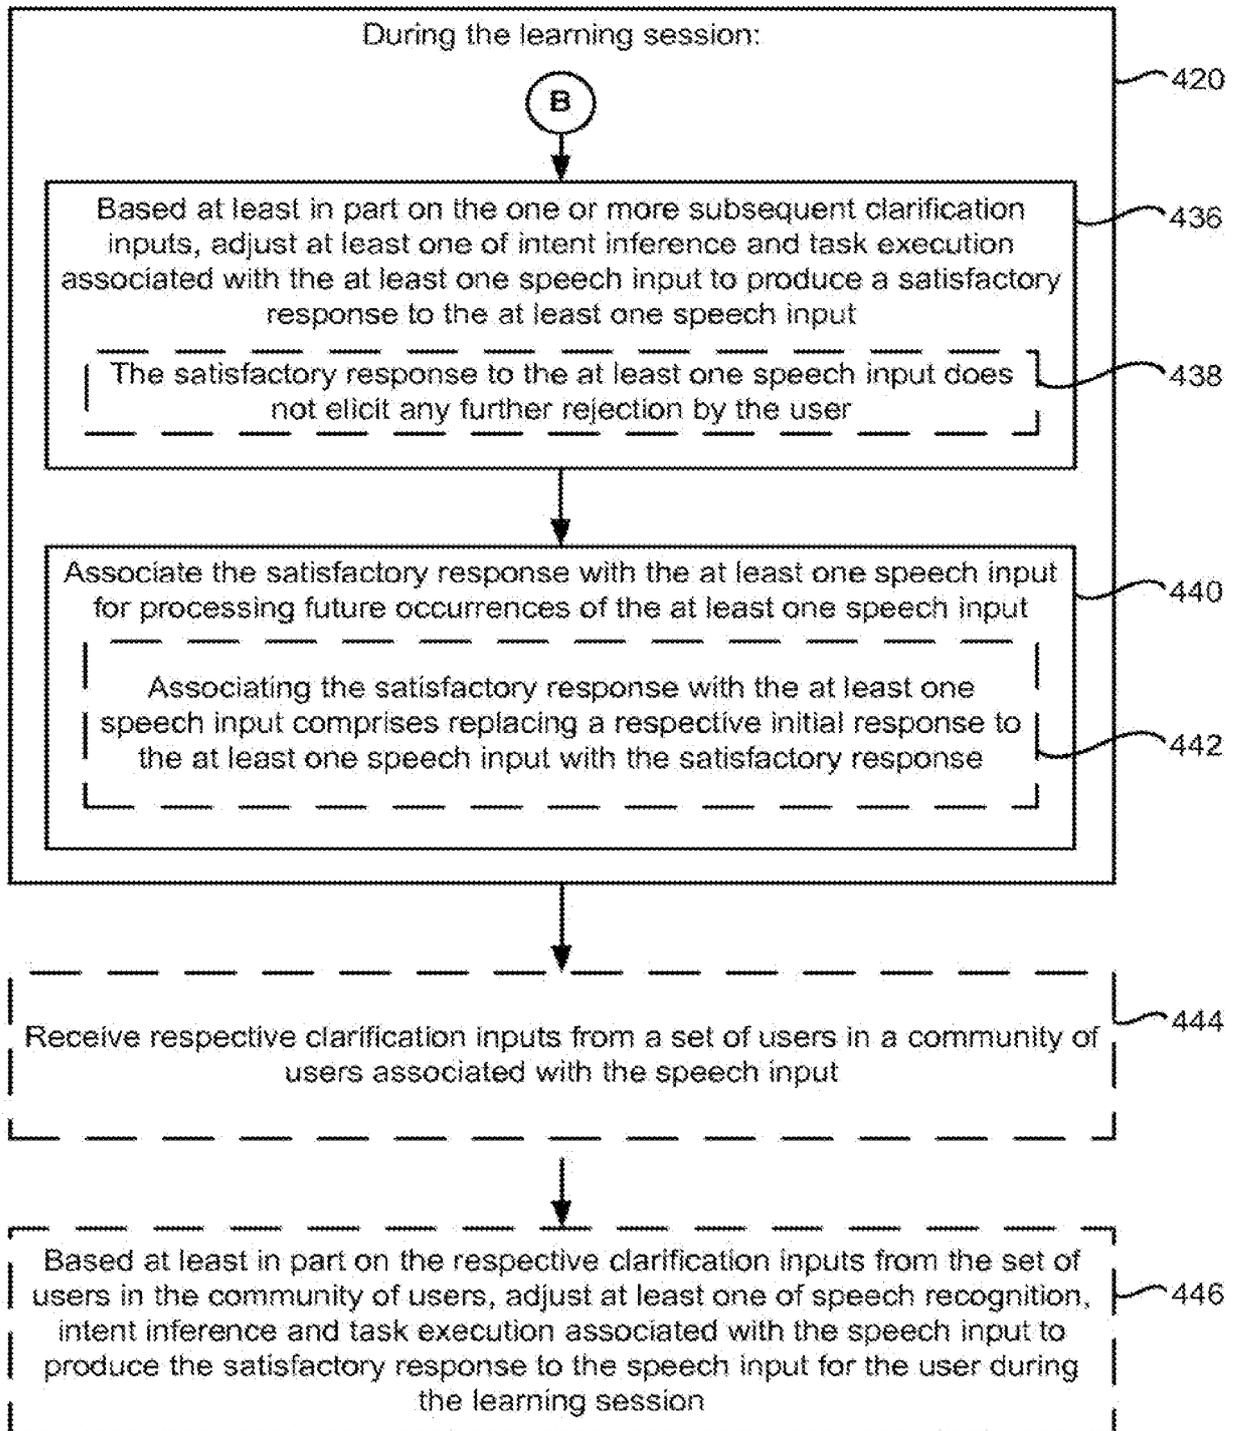


Figure 4C

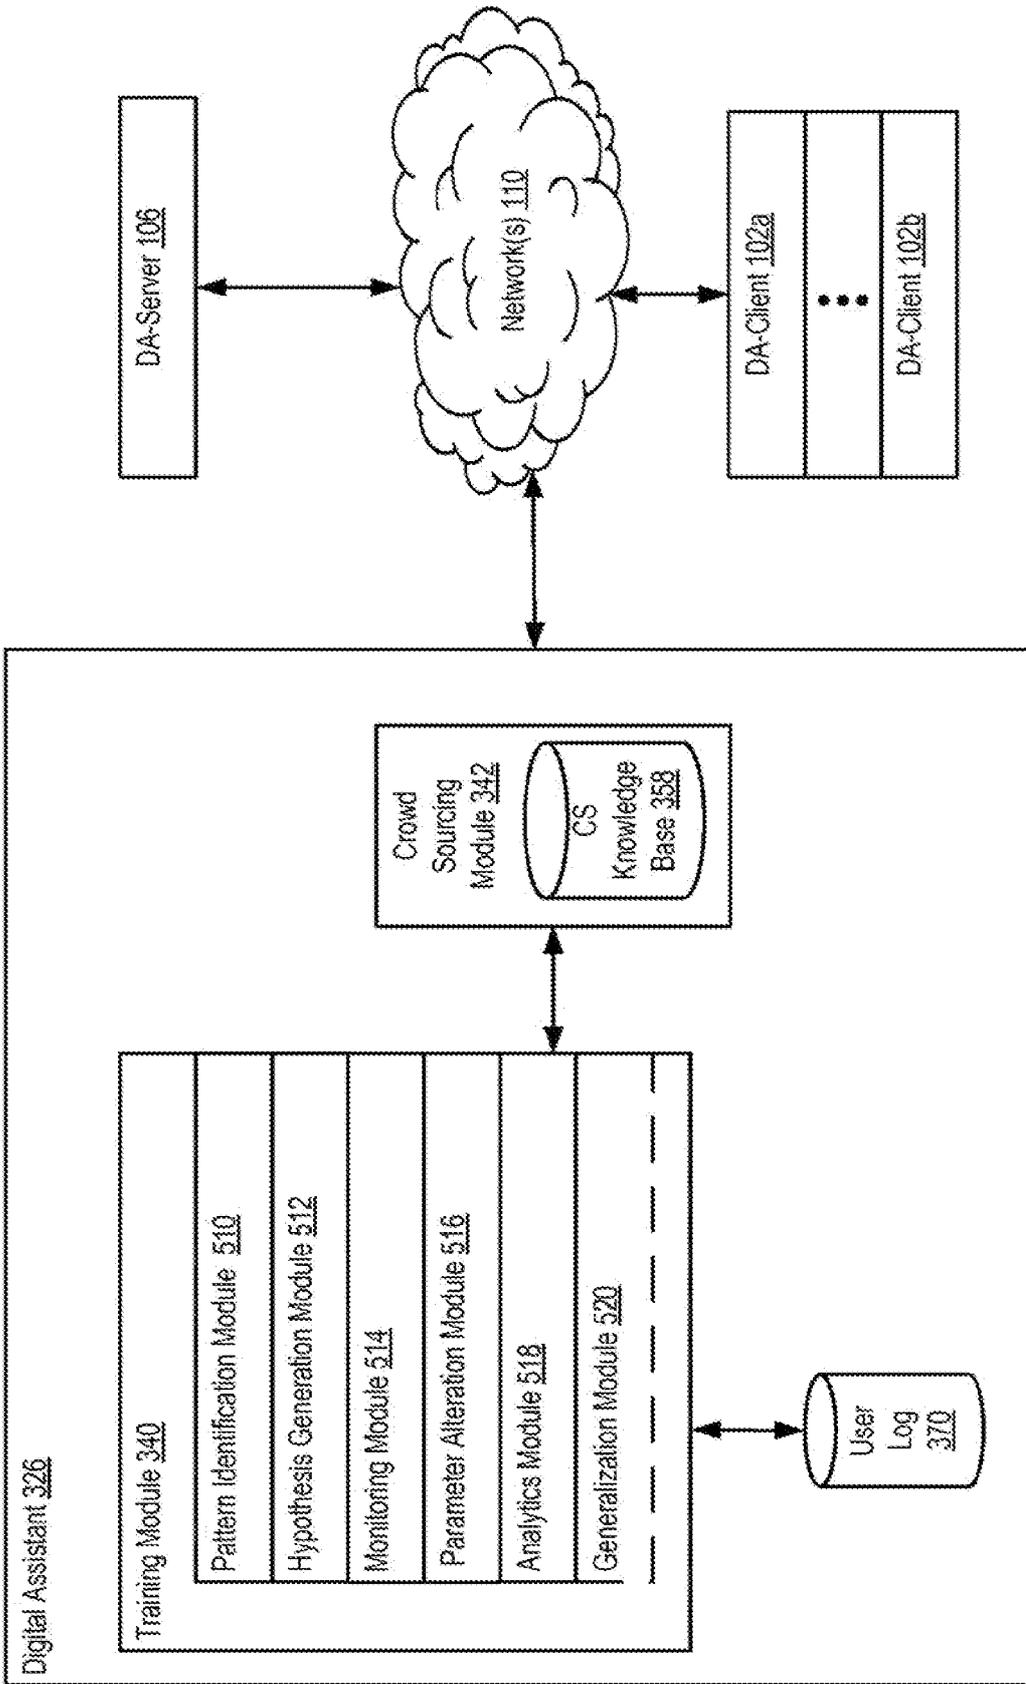


Figure 5

600

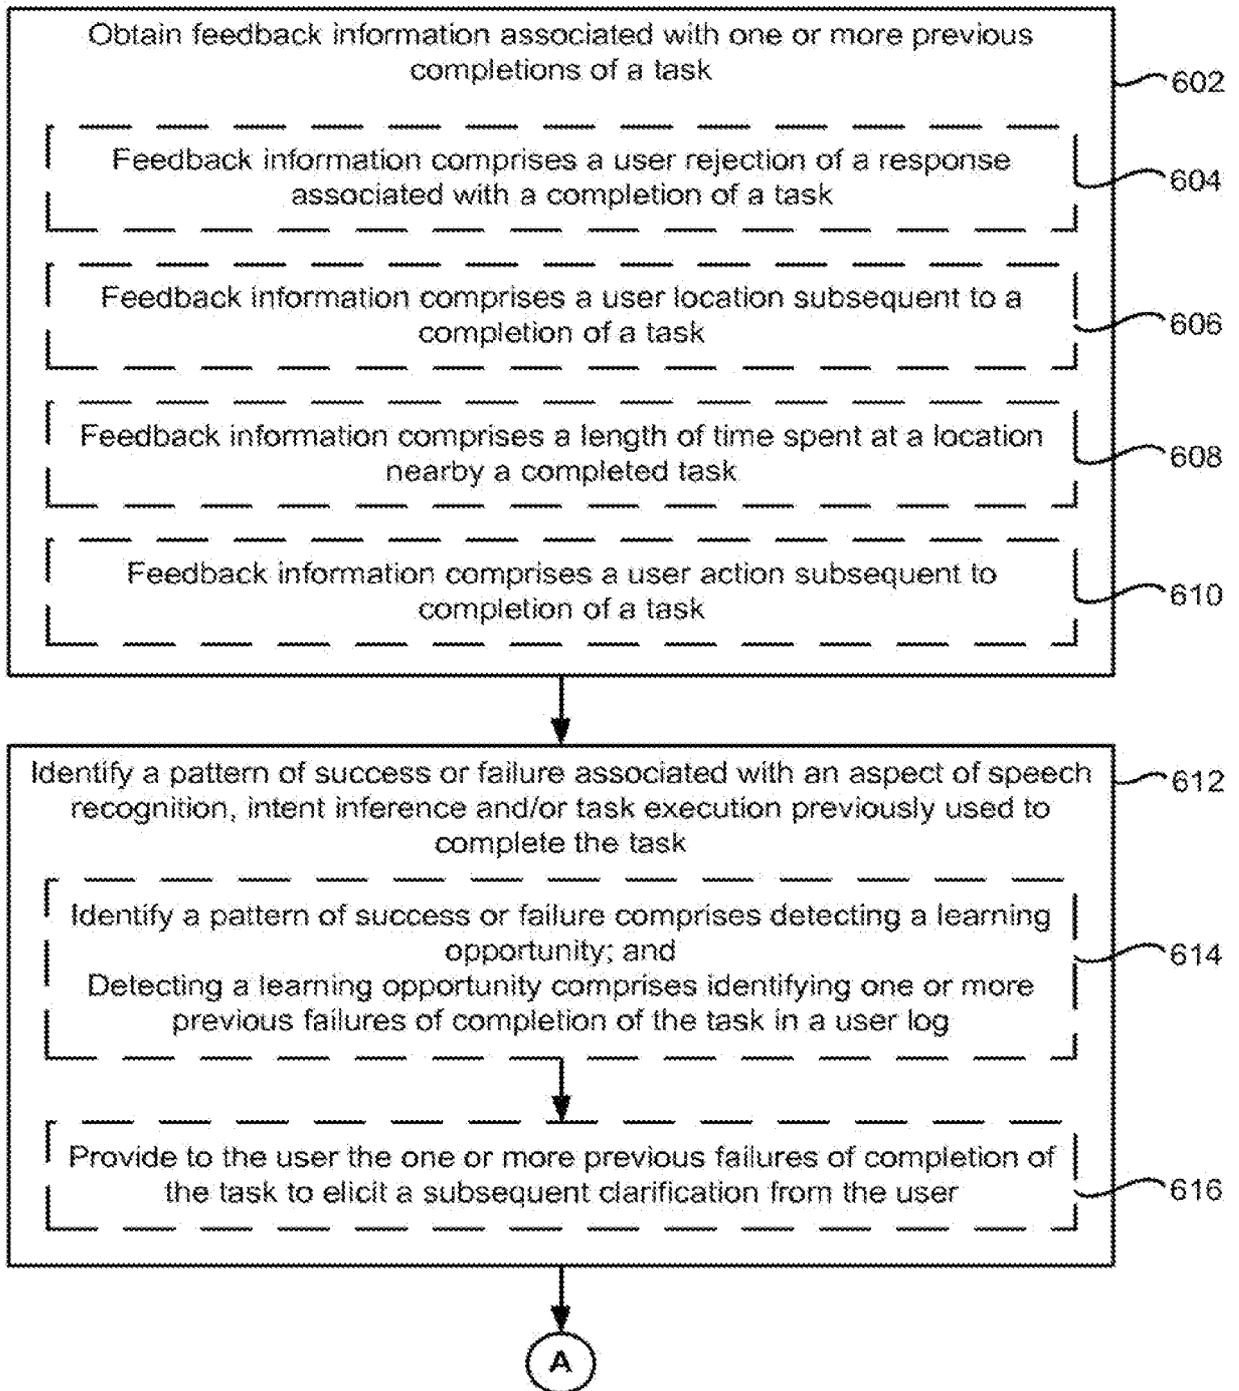


Figure 6A

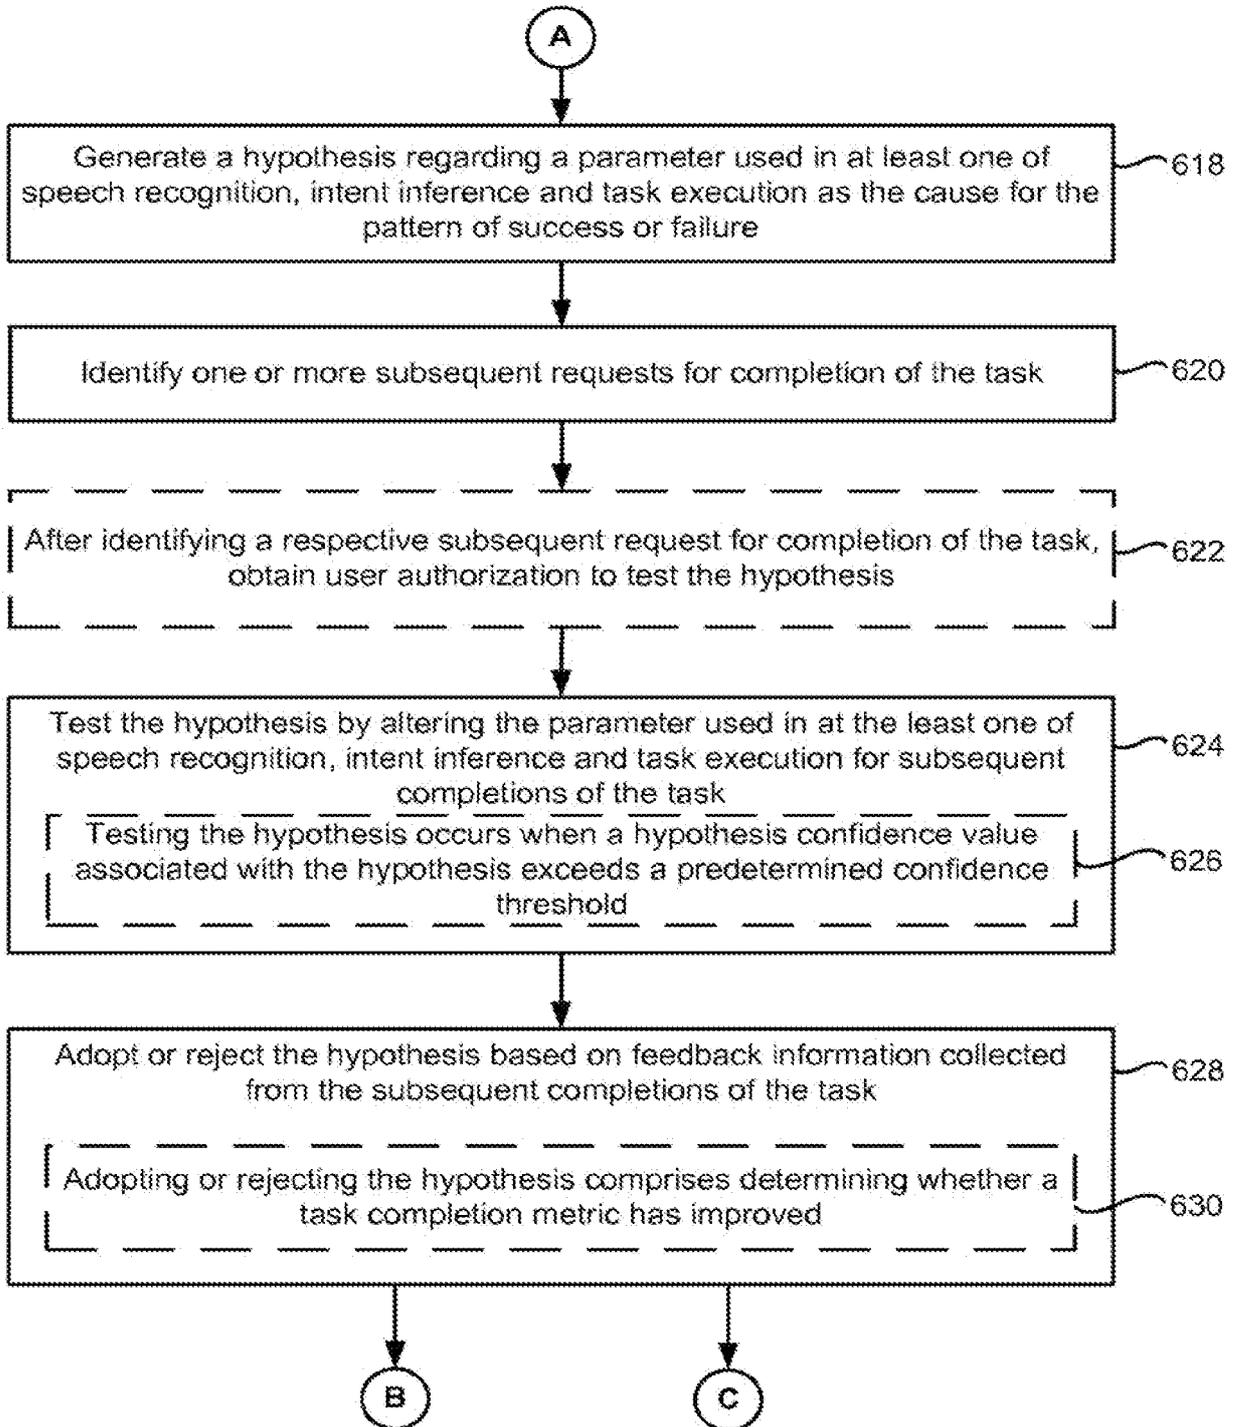


Figure 6B

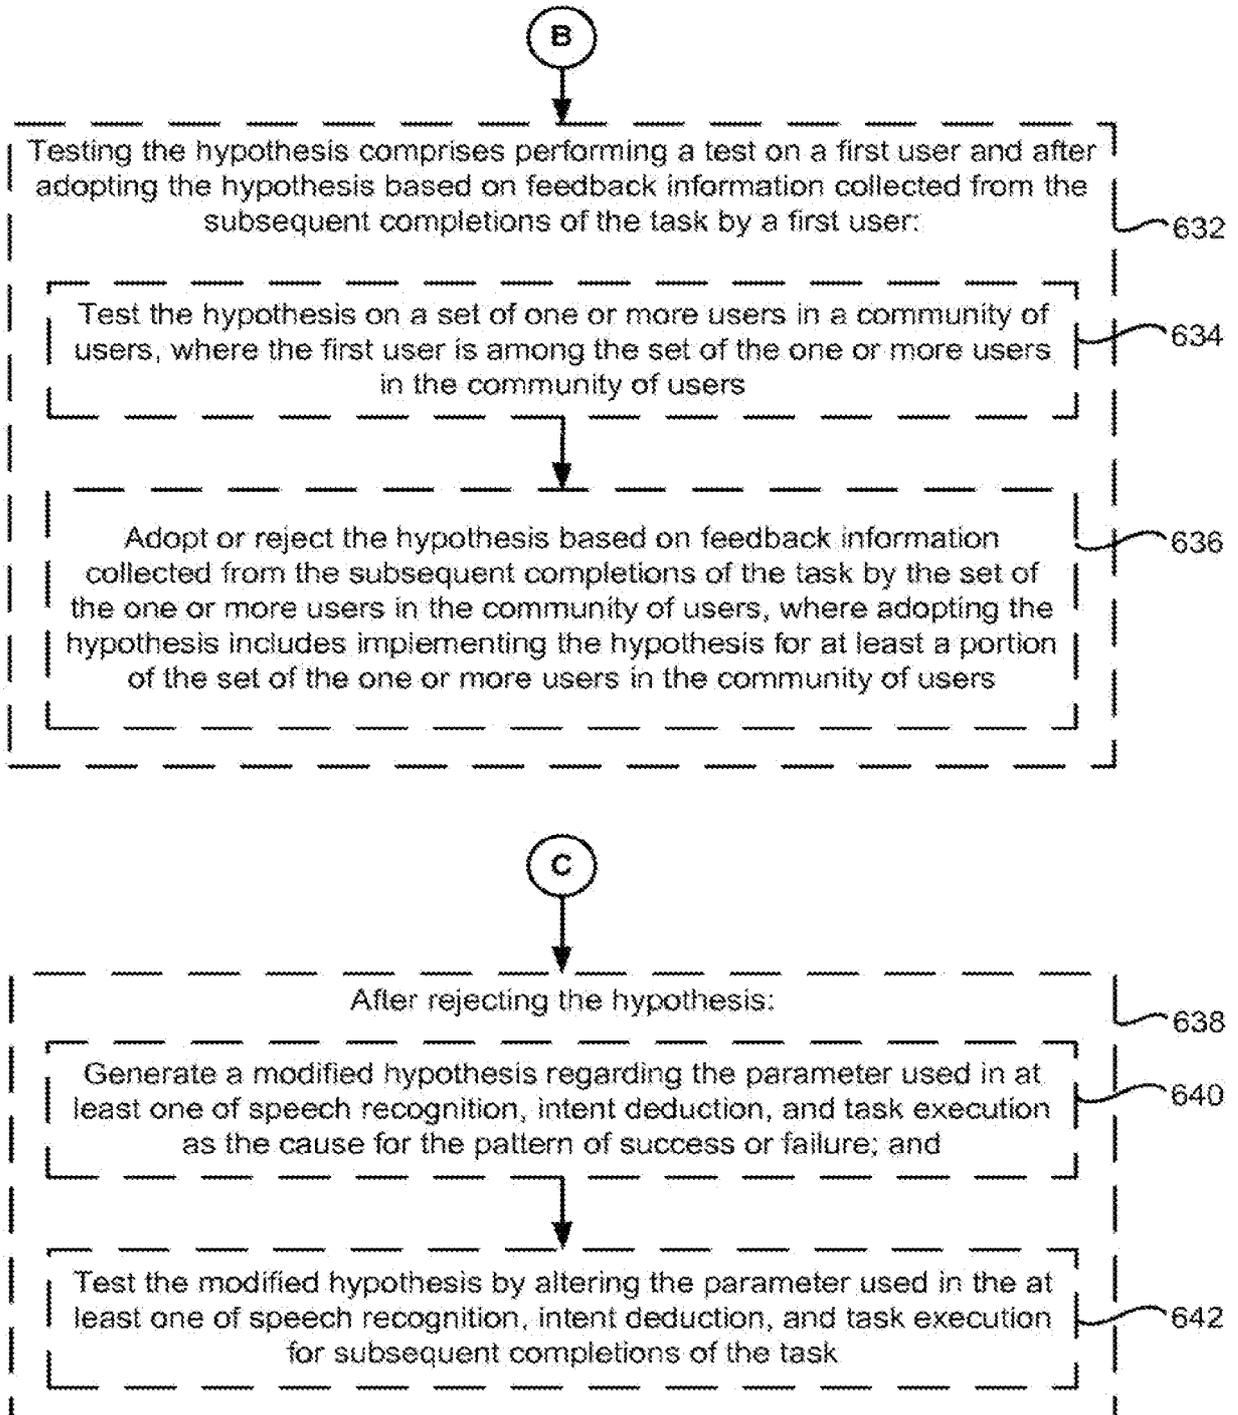


Figure 6C

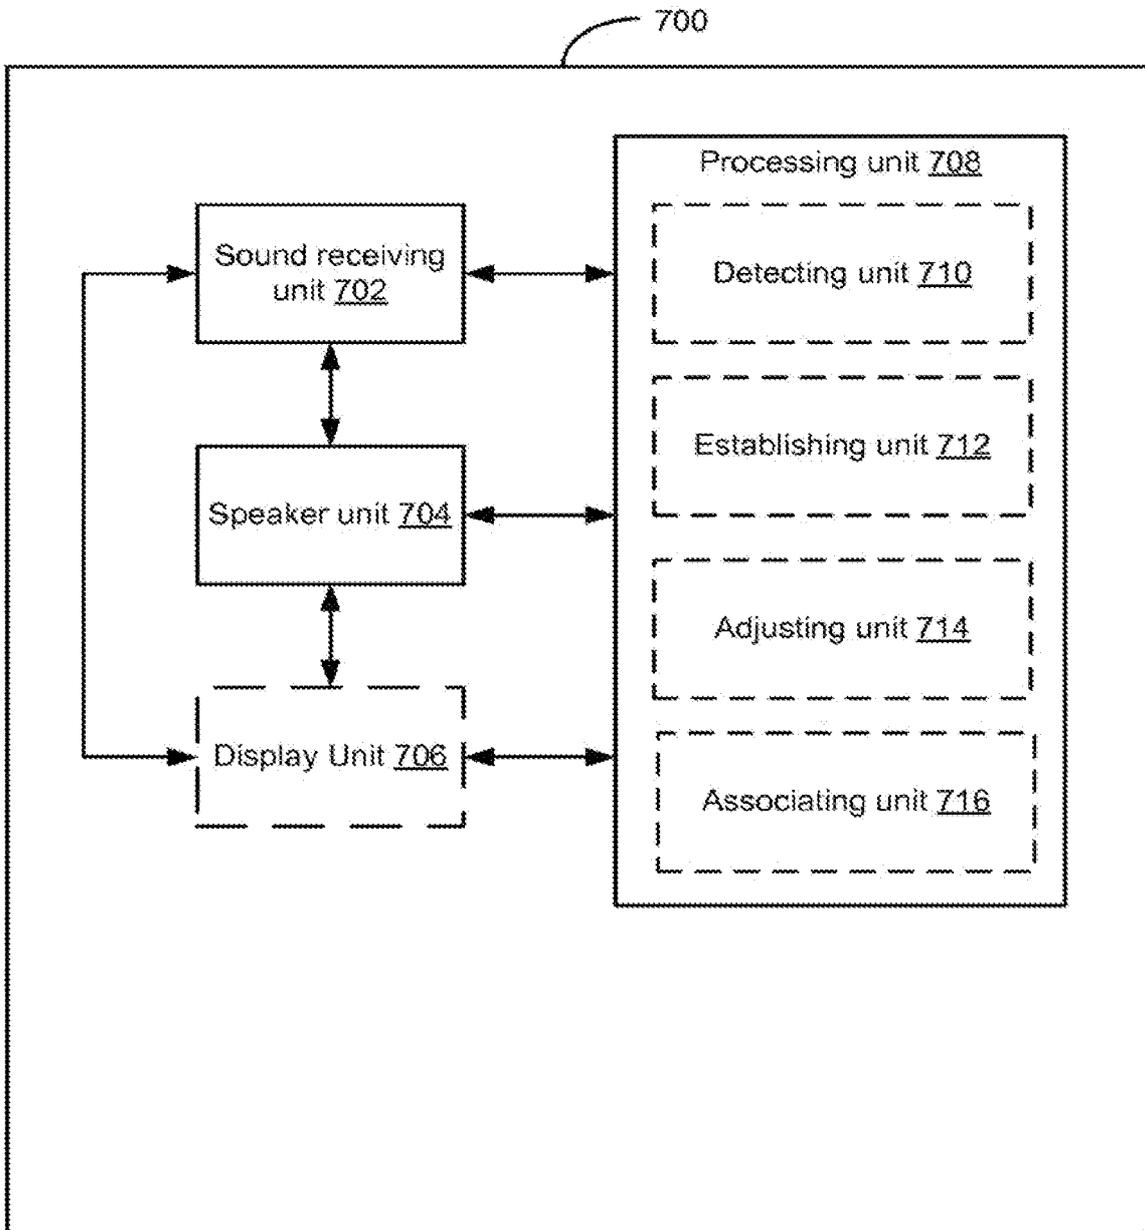


Figure 7

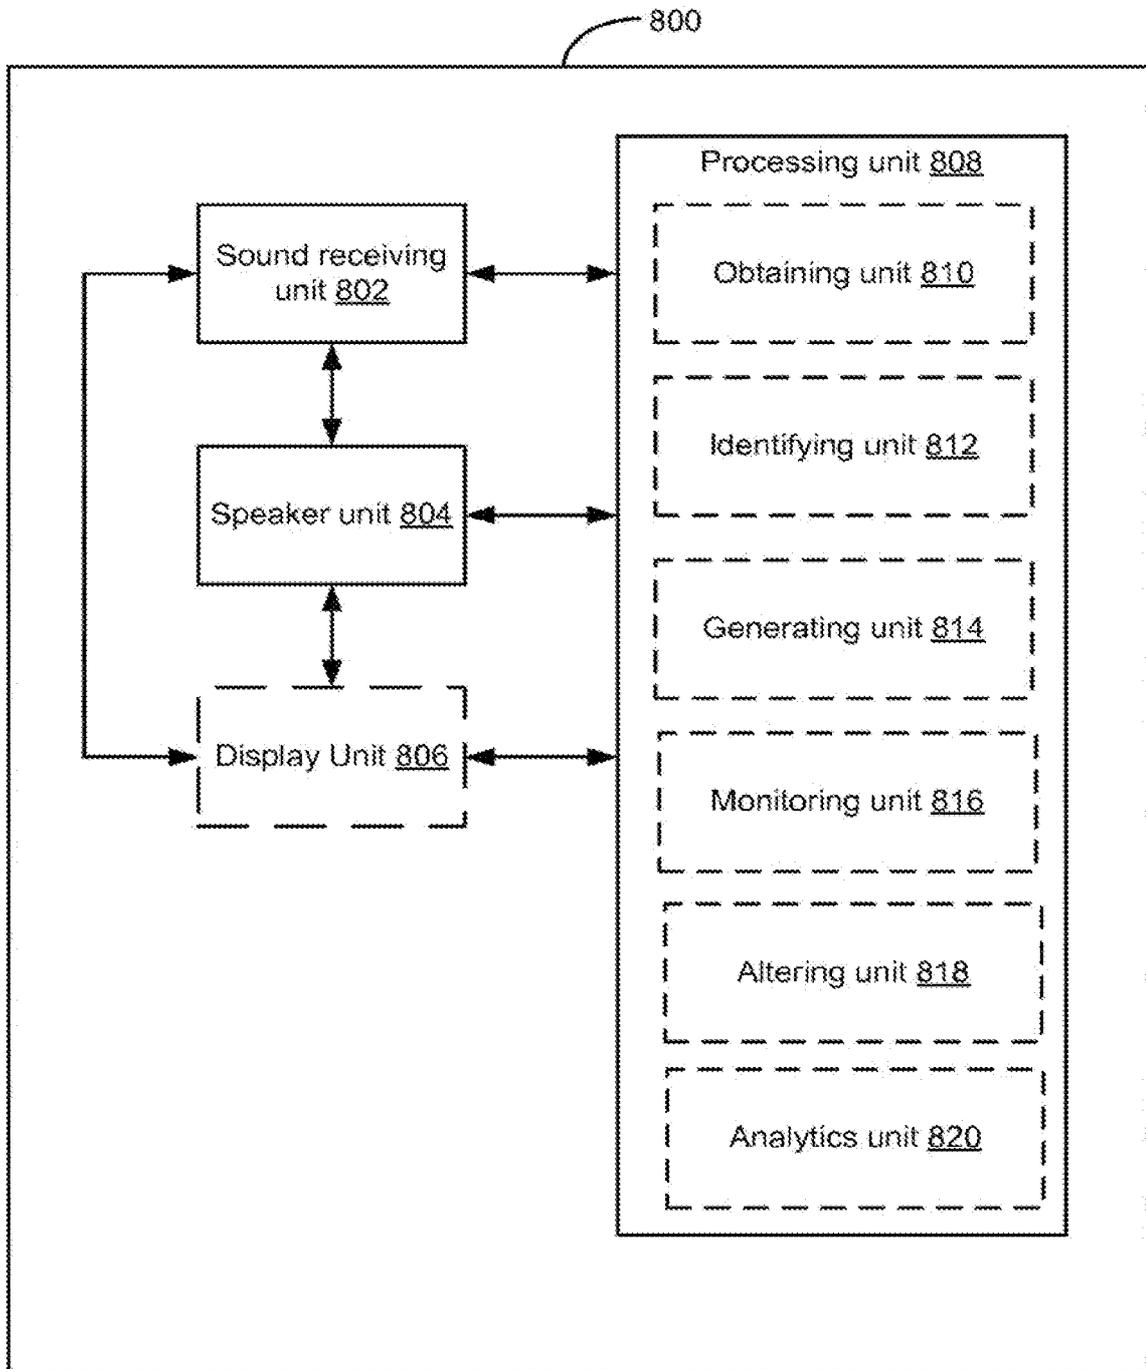


Figure 8