



(19) **United States**

(12) **Patent Application Publication**  
**GARCIA et al.**

(10) **Pub. No.: US 2023/0335132 A1**

(43) **Pub. Date: Oct. 19, 2023**

(54) **NATURAL ASSISTANT INTERACTION**

**G06F 3/16** (2006.01)

**G06F 40/30** (2006.01)

**G10L 15/26** (2006.01)

(71) Applicant: **Apple Inc.**, Cupertino, CA (US)

(72) Inventors: **Juan Carlos GARCIA**, San Francisco, CA (US); **Paul S. MCCARTHY**, Traversse City, MI (US); **Kurt PIERSOL**, San Jose, CA (US)

(52) **U.S. Cl.**

CPC ..... **G10L 15/22** (2013.01); **G10L 25/78** (2013.01); **G10L 15/08** (2013.01); **G06F 3/167** (2013.01); **G06F 40/30** (2020.01); **G10L 15/26** (2013.01); **G10L 2025/783** (2013.01); **G10L 2015/223** (2013.01); **G10L 2015/088** (2013.01); **G10L 25/87** (2013.01)

(21) Appl. No.: **18/212,654**

(22) Filed: **Jun. 21, 2023**

(57) **ABSTRACT**

**Related U.S. Application Data**

(63) Continuation of application No. 17/066,228, filed on Oct. 8, 2020, now Pat. No. 11,710,482, which is a continuation of application No. 16/019,331, filed on Jun. 26, 2018, now Pat. No. 10,818,288.

(60) Provisional application No. 62/648,084, filed on Mar. 26, 2018.

Systems and processes for operating a virtual assistant to provide natural assistant interaction are provided. In accordance with one or more examples, a method includes, at an electronic device with one or more processors and memory: receiving a first audio stream including one or more utterances; determining whether the first audio stream includes a lexical trigger; generating one or more candidate text representations of the one or more utterances; determining whether at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant. If at least one candidate text representation is to be disregarded, one or more candidate intents are generated based on candidate text representations of the one or more candidate text representations other than the to be disregarded at least one candidate text representation.

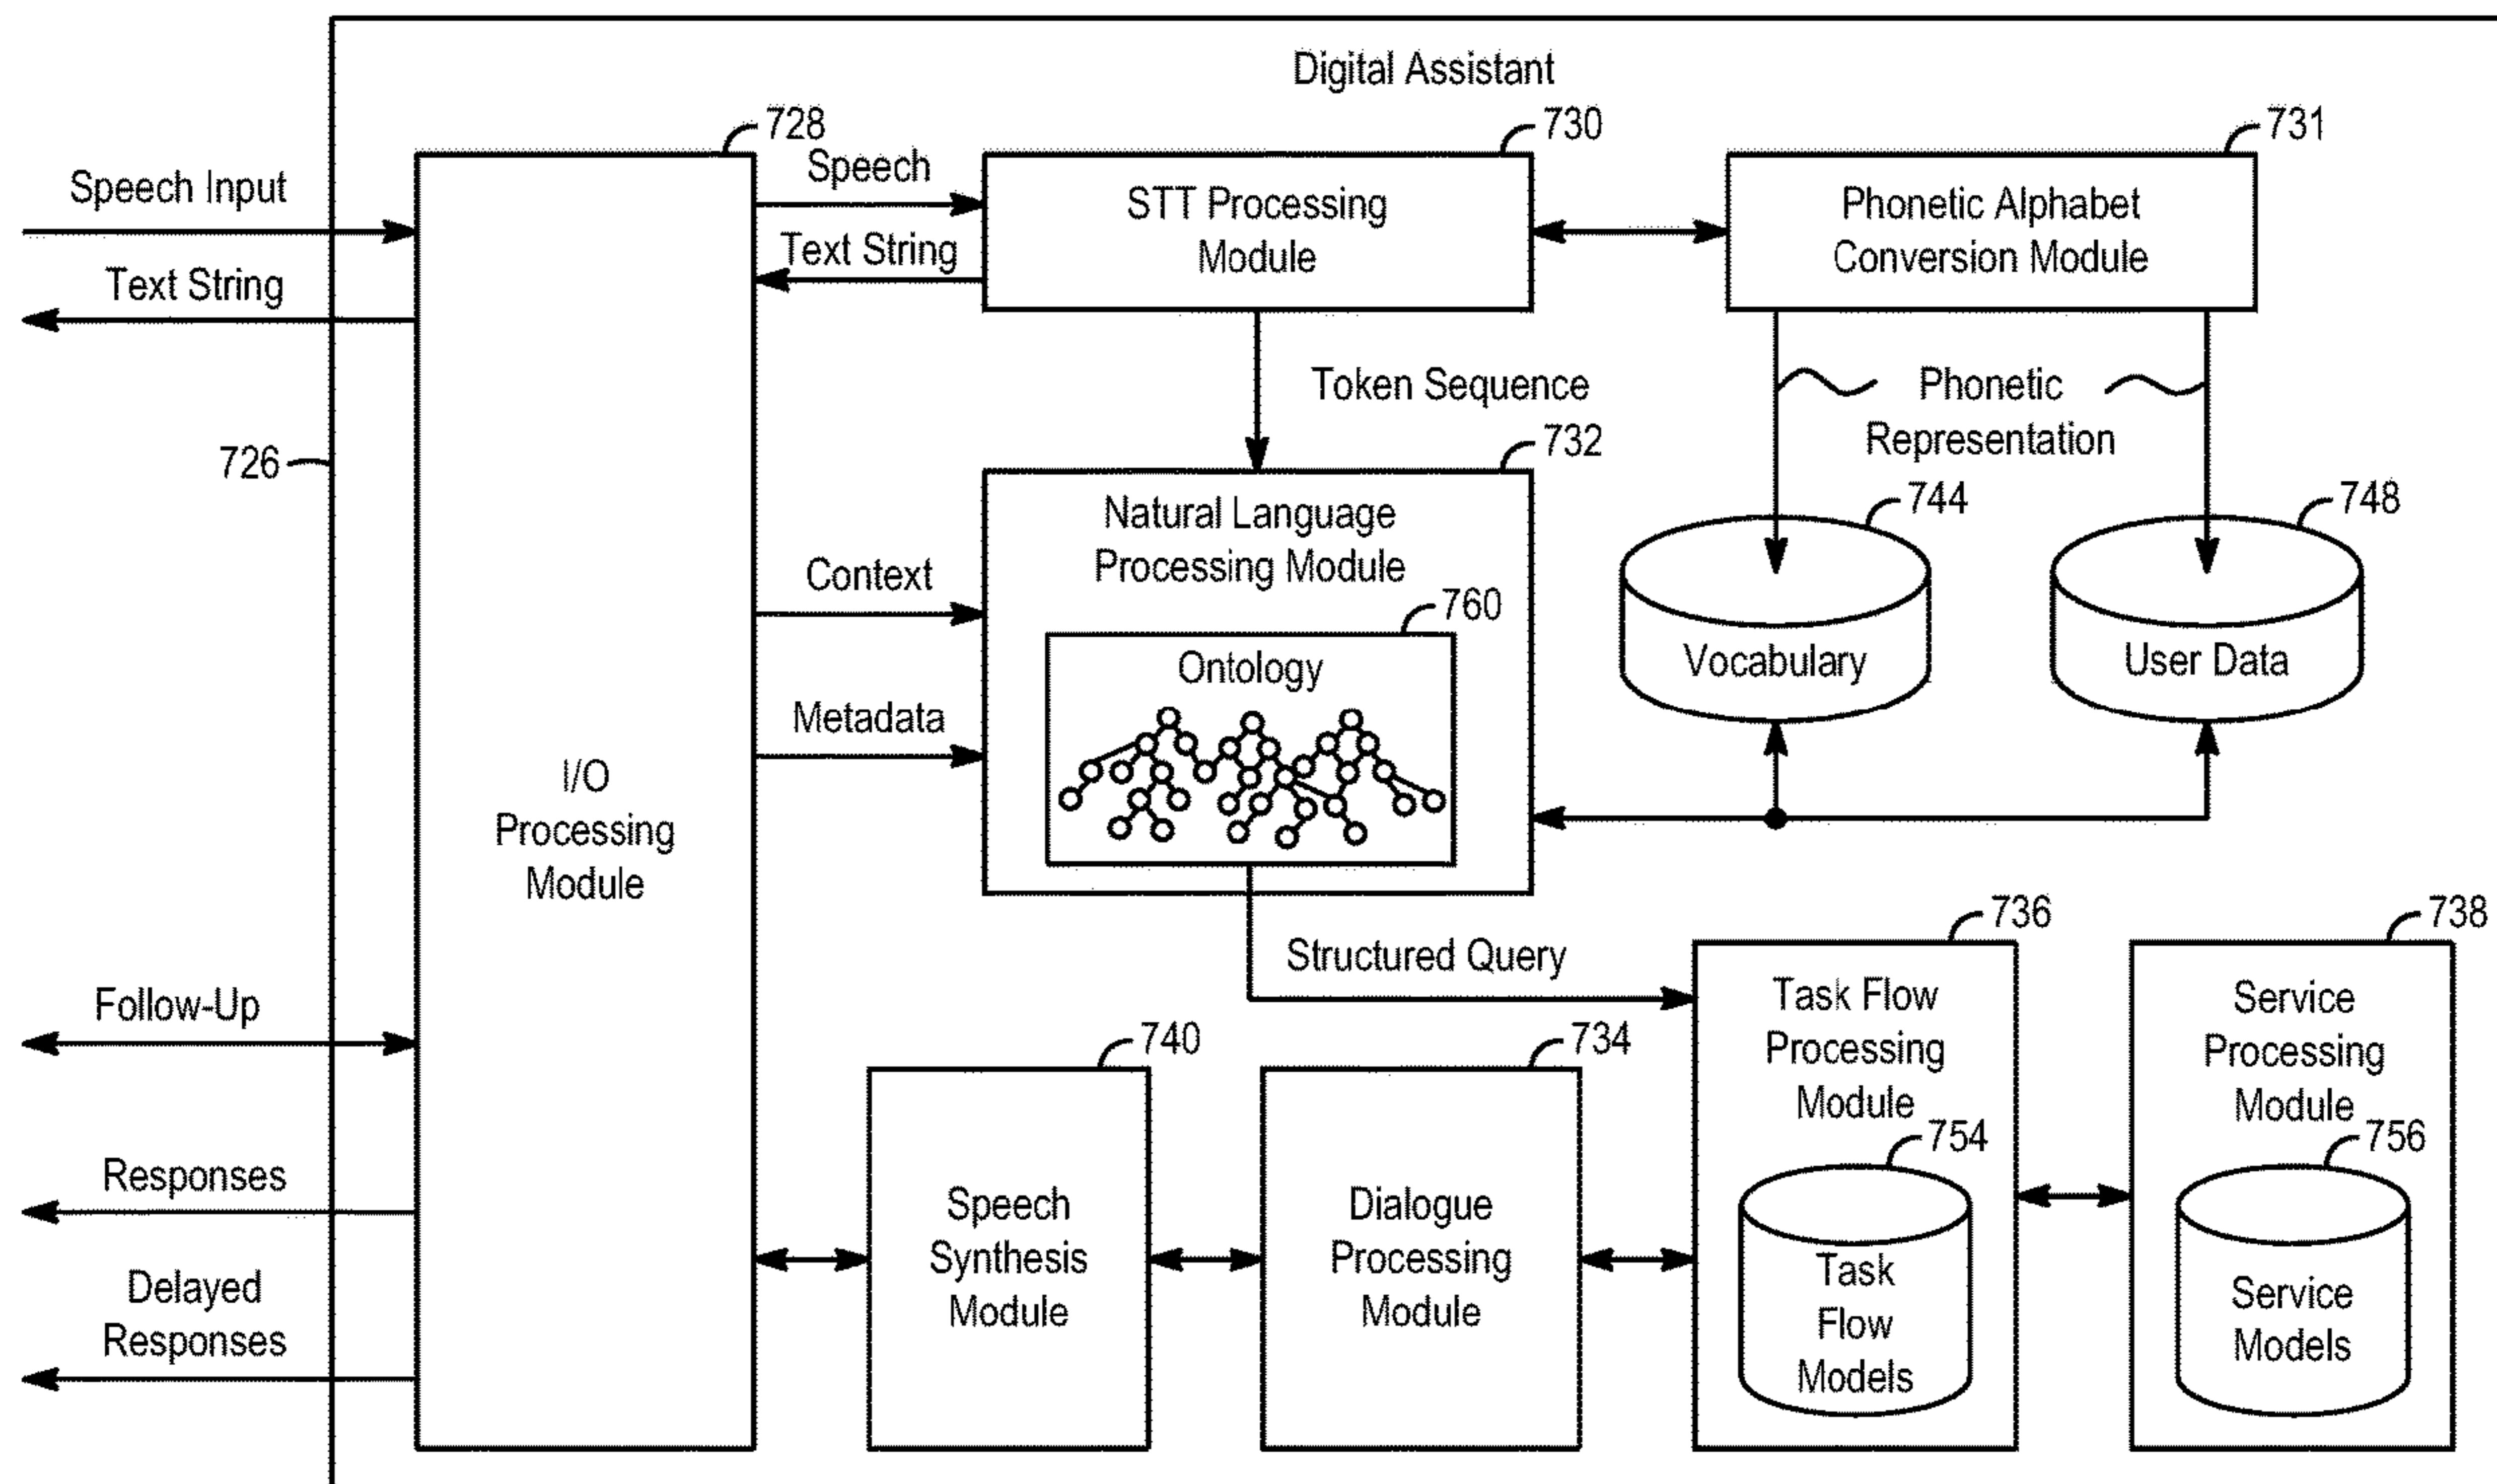
**Publication Classification**

(51) **Int. Cl.**

**G10L 15/22** (2006.01)

**G10L 25/78** (2006.01)

**G10L 15/08** (2006.01)






FIG. 1




FIG. 2A

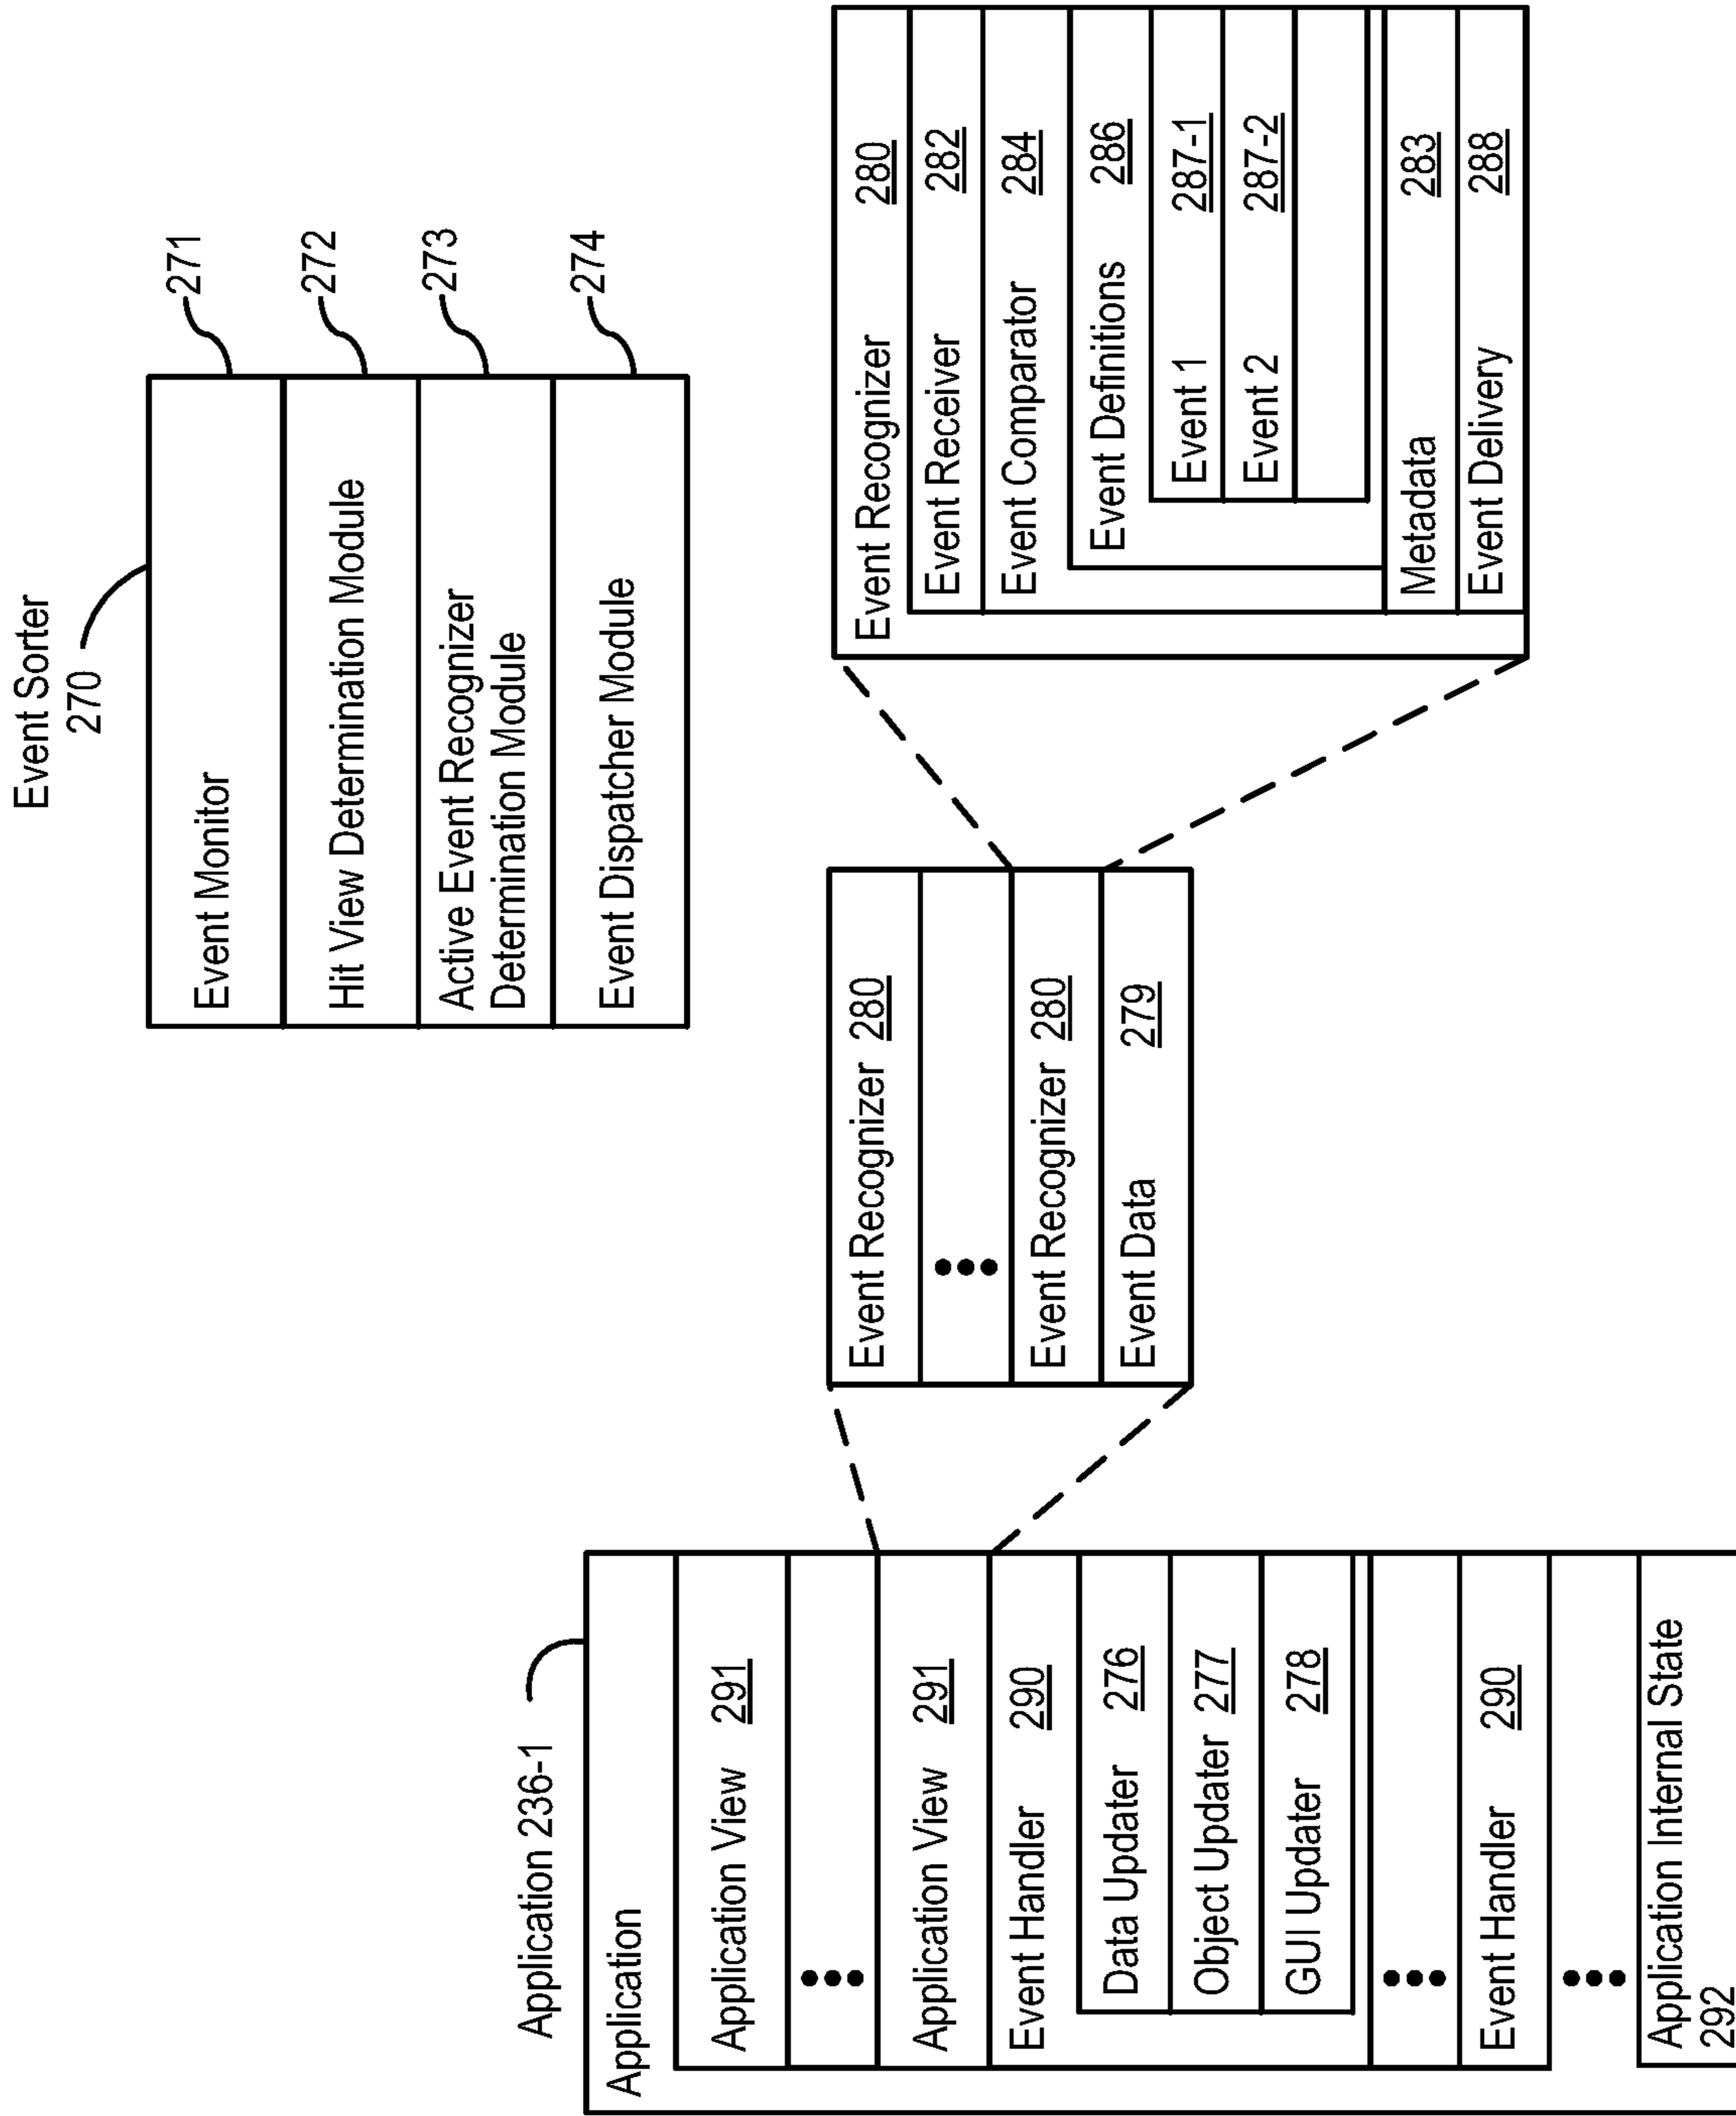


FIG. 2B




FIG. 3

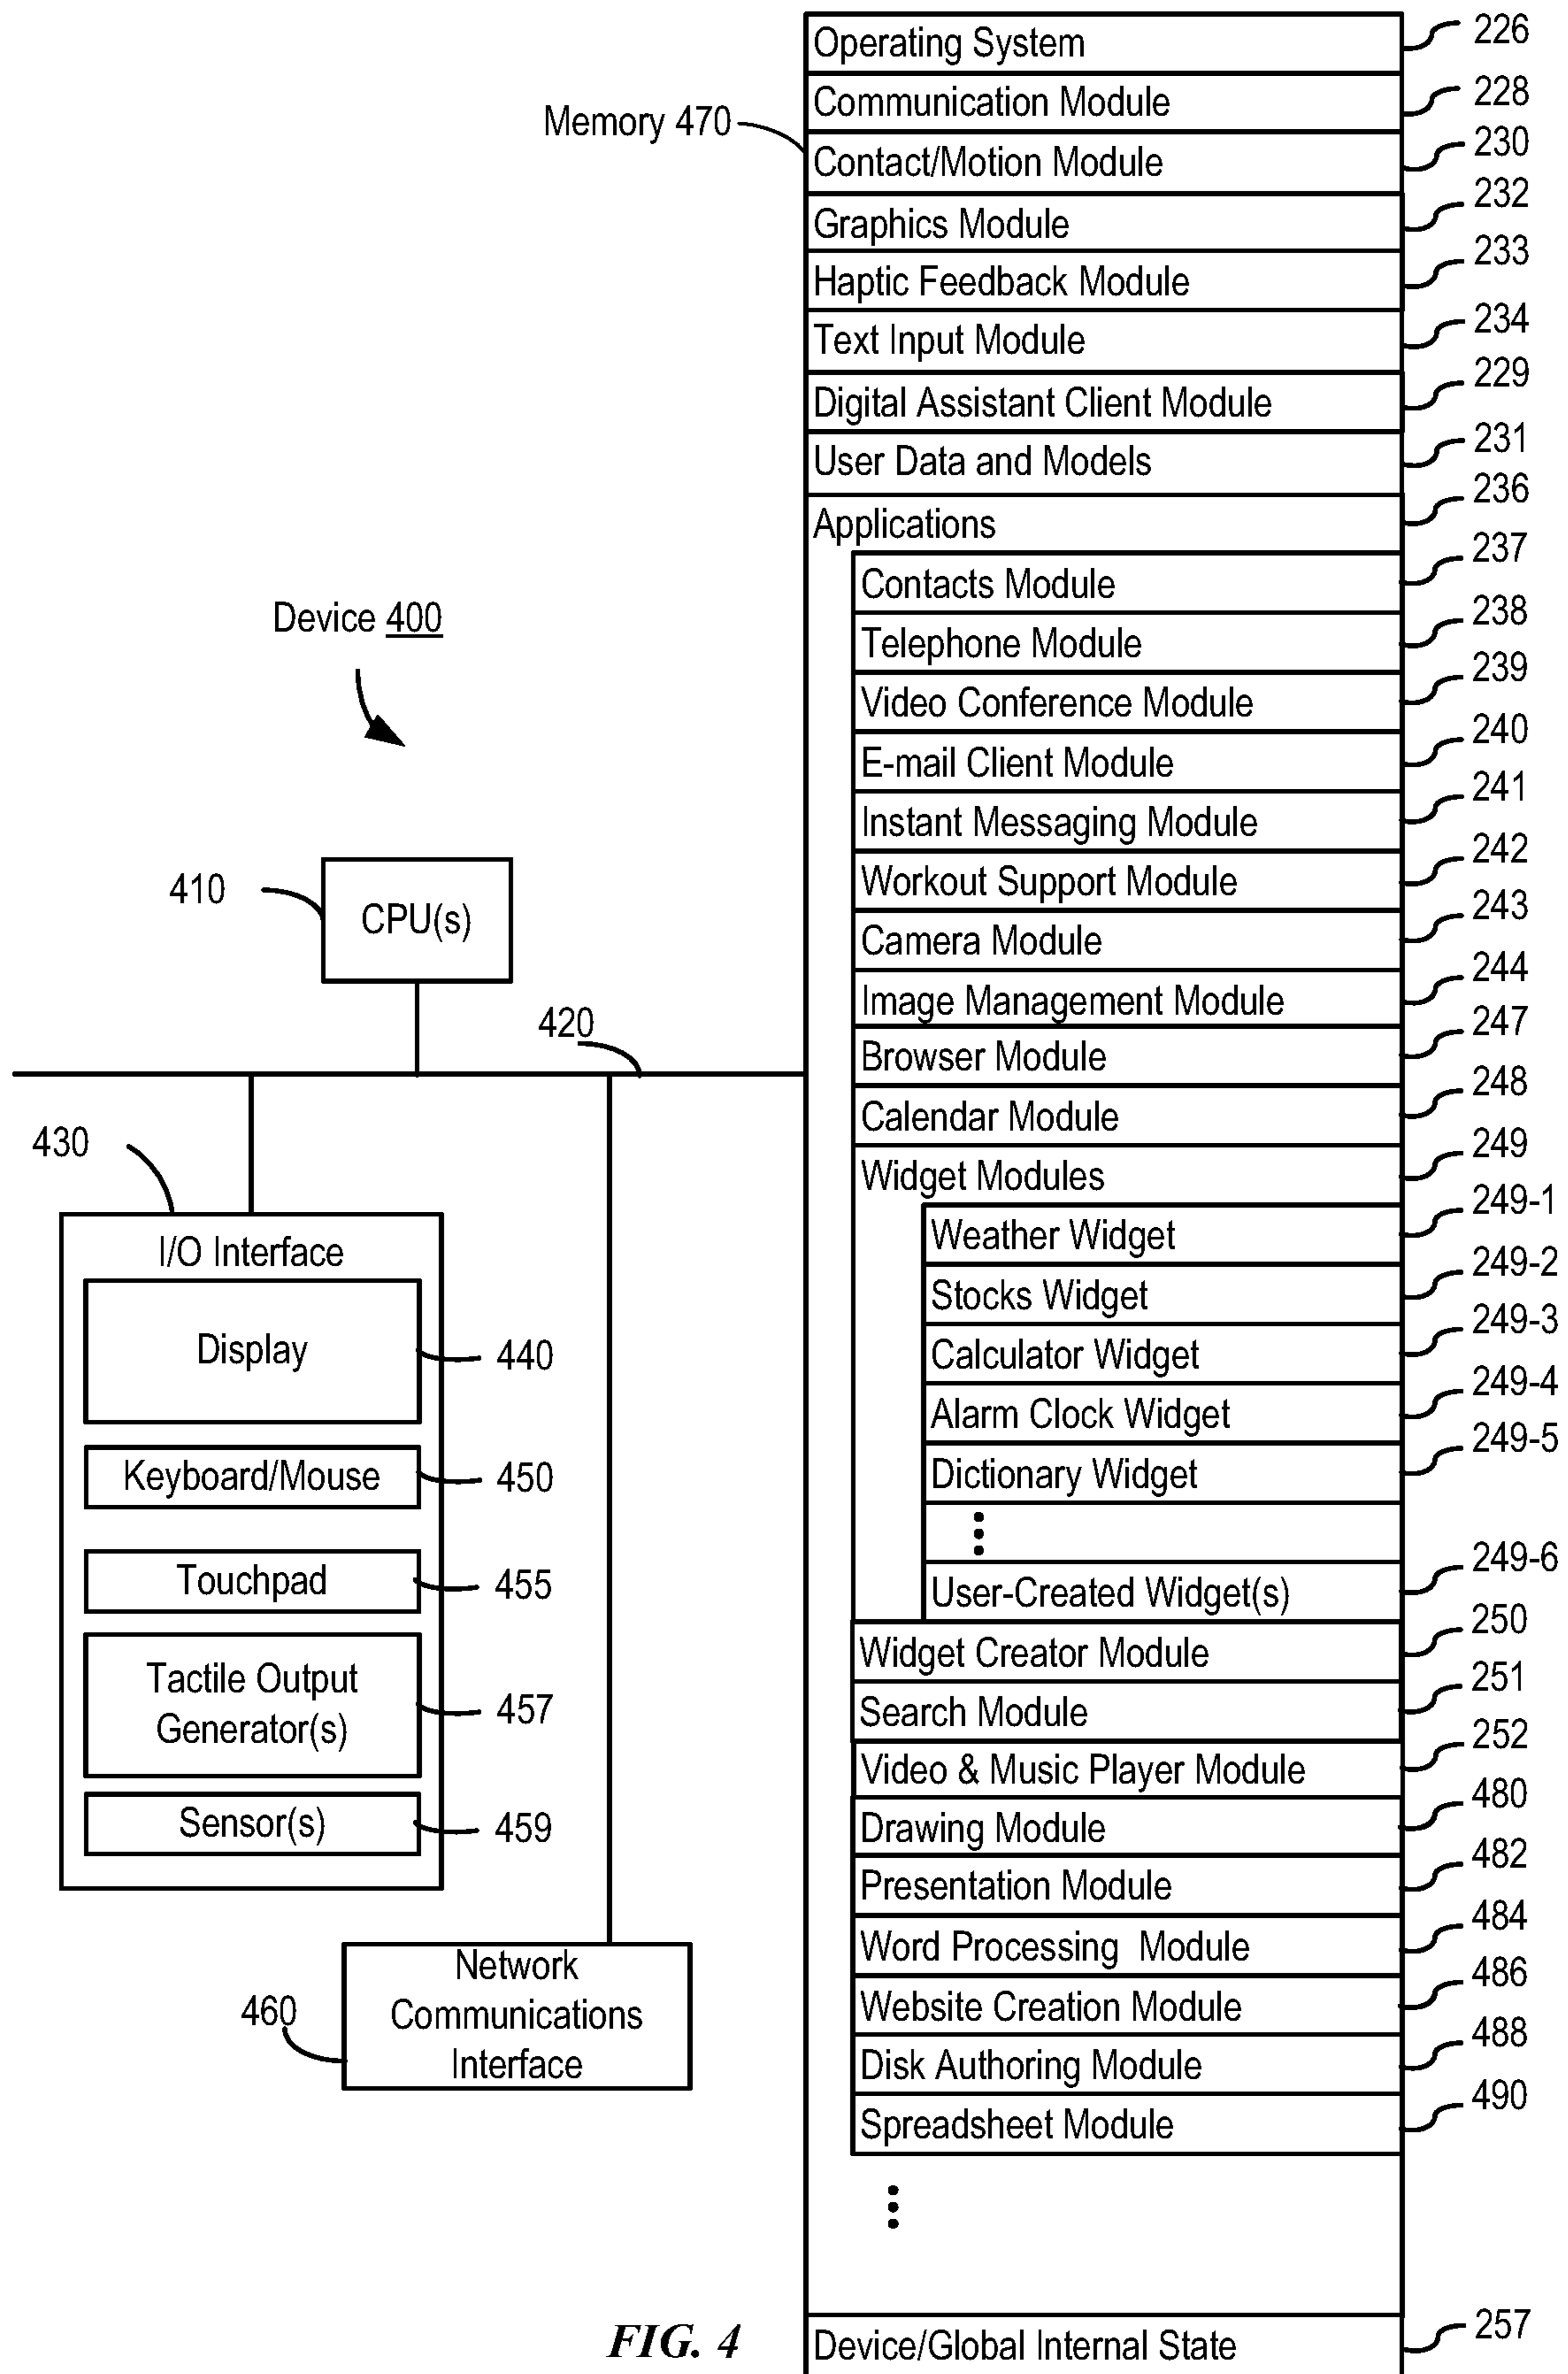


FIG. 4




FIG. 5A






FIG. 5B





**FIG. 6A**




FIG. 6B

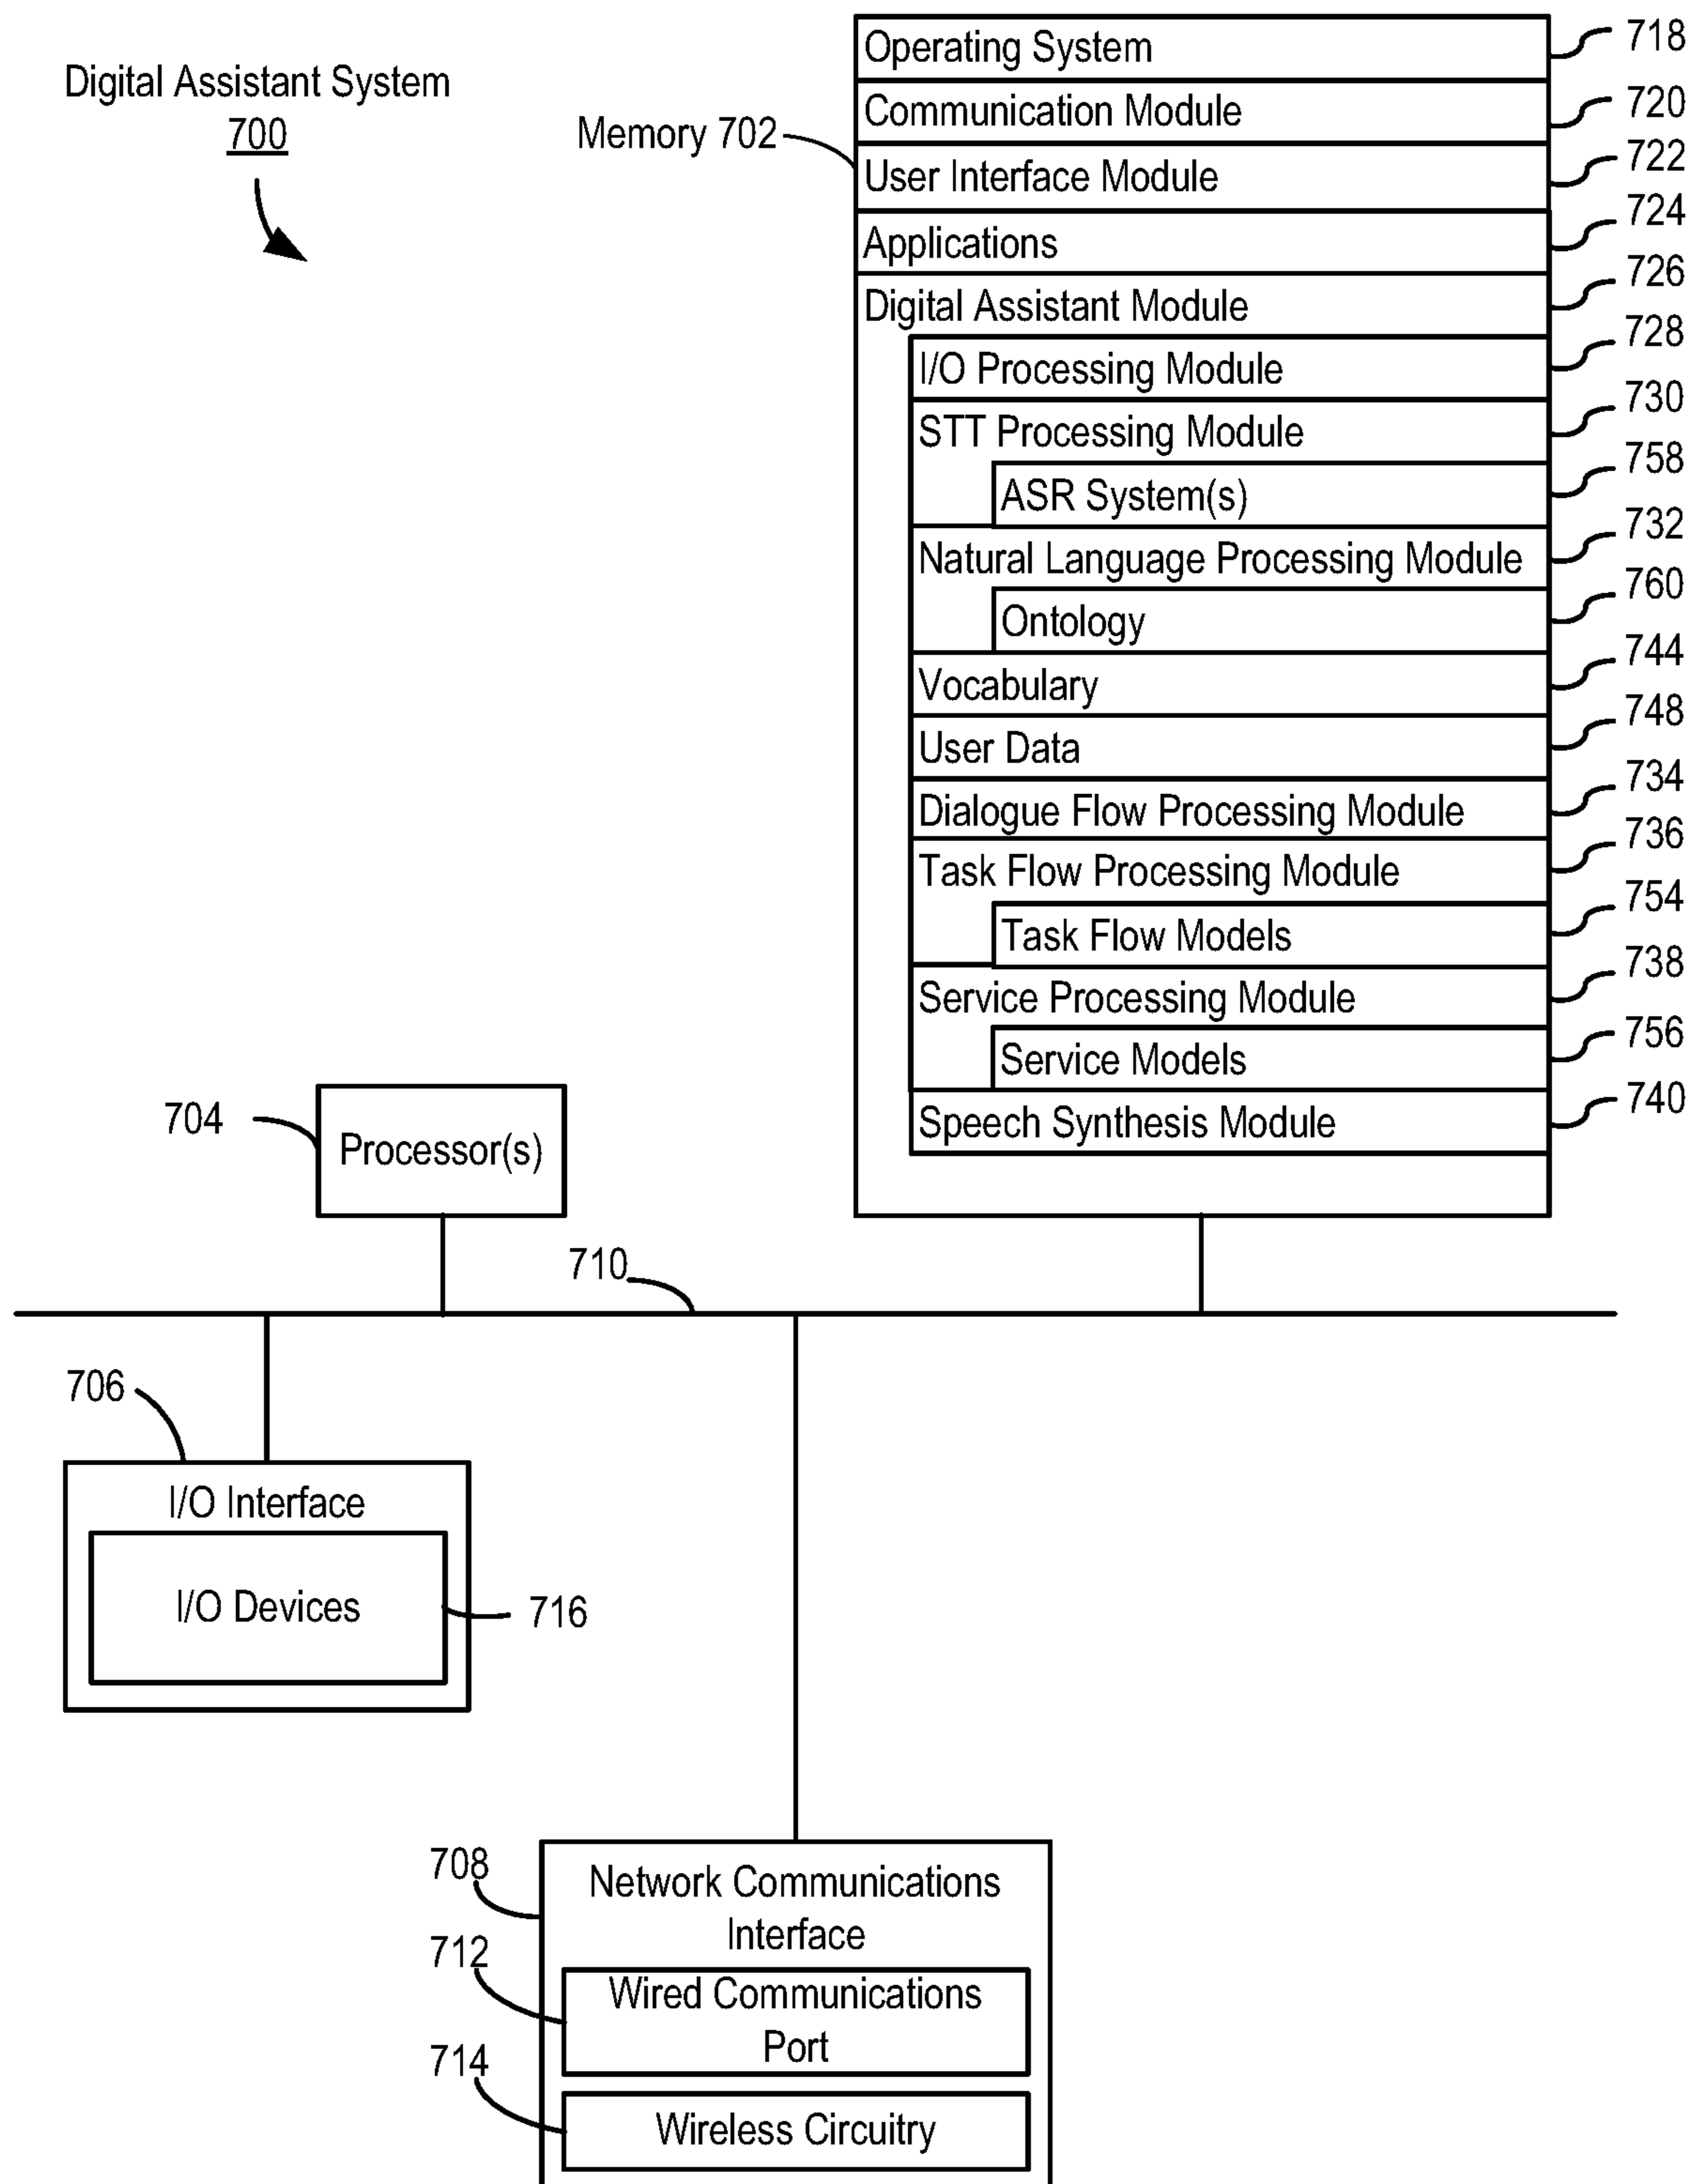


FIG. 7A

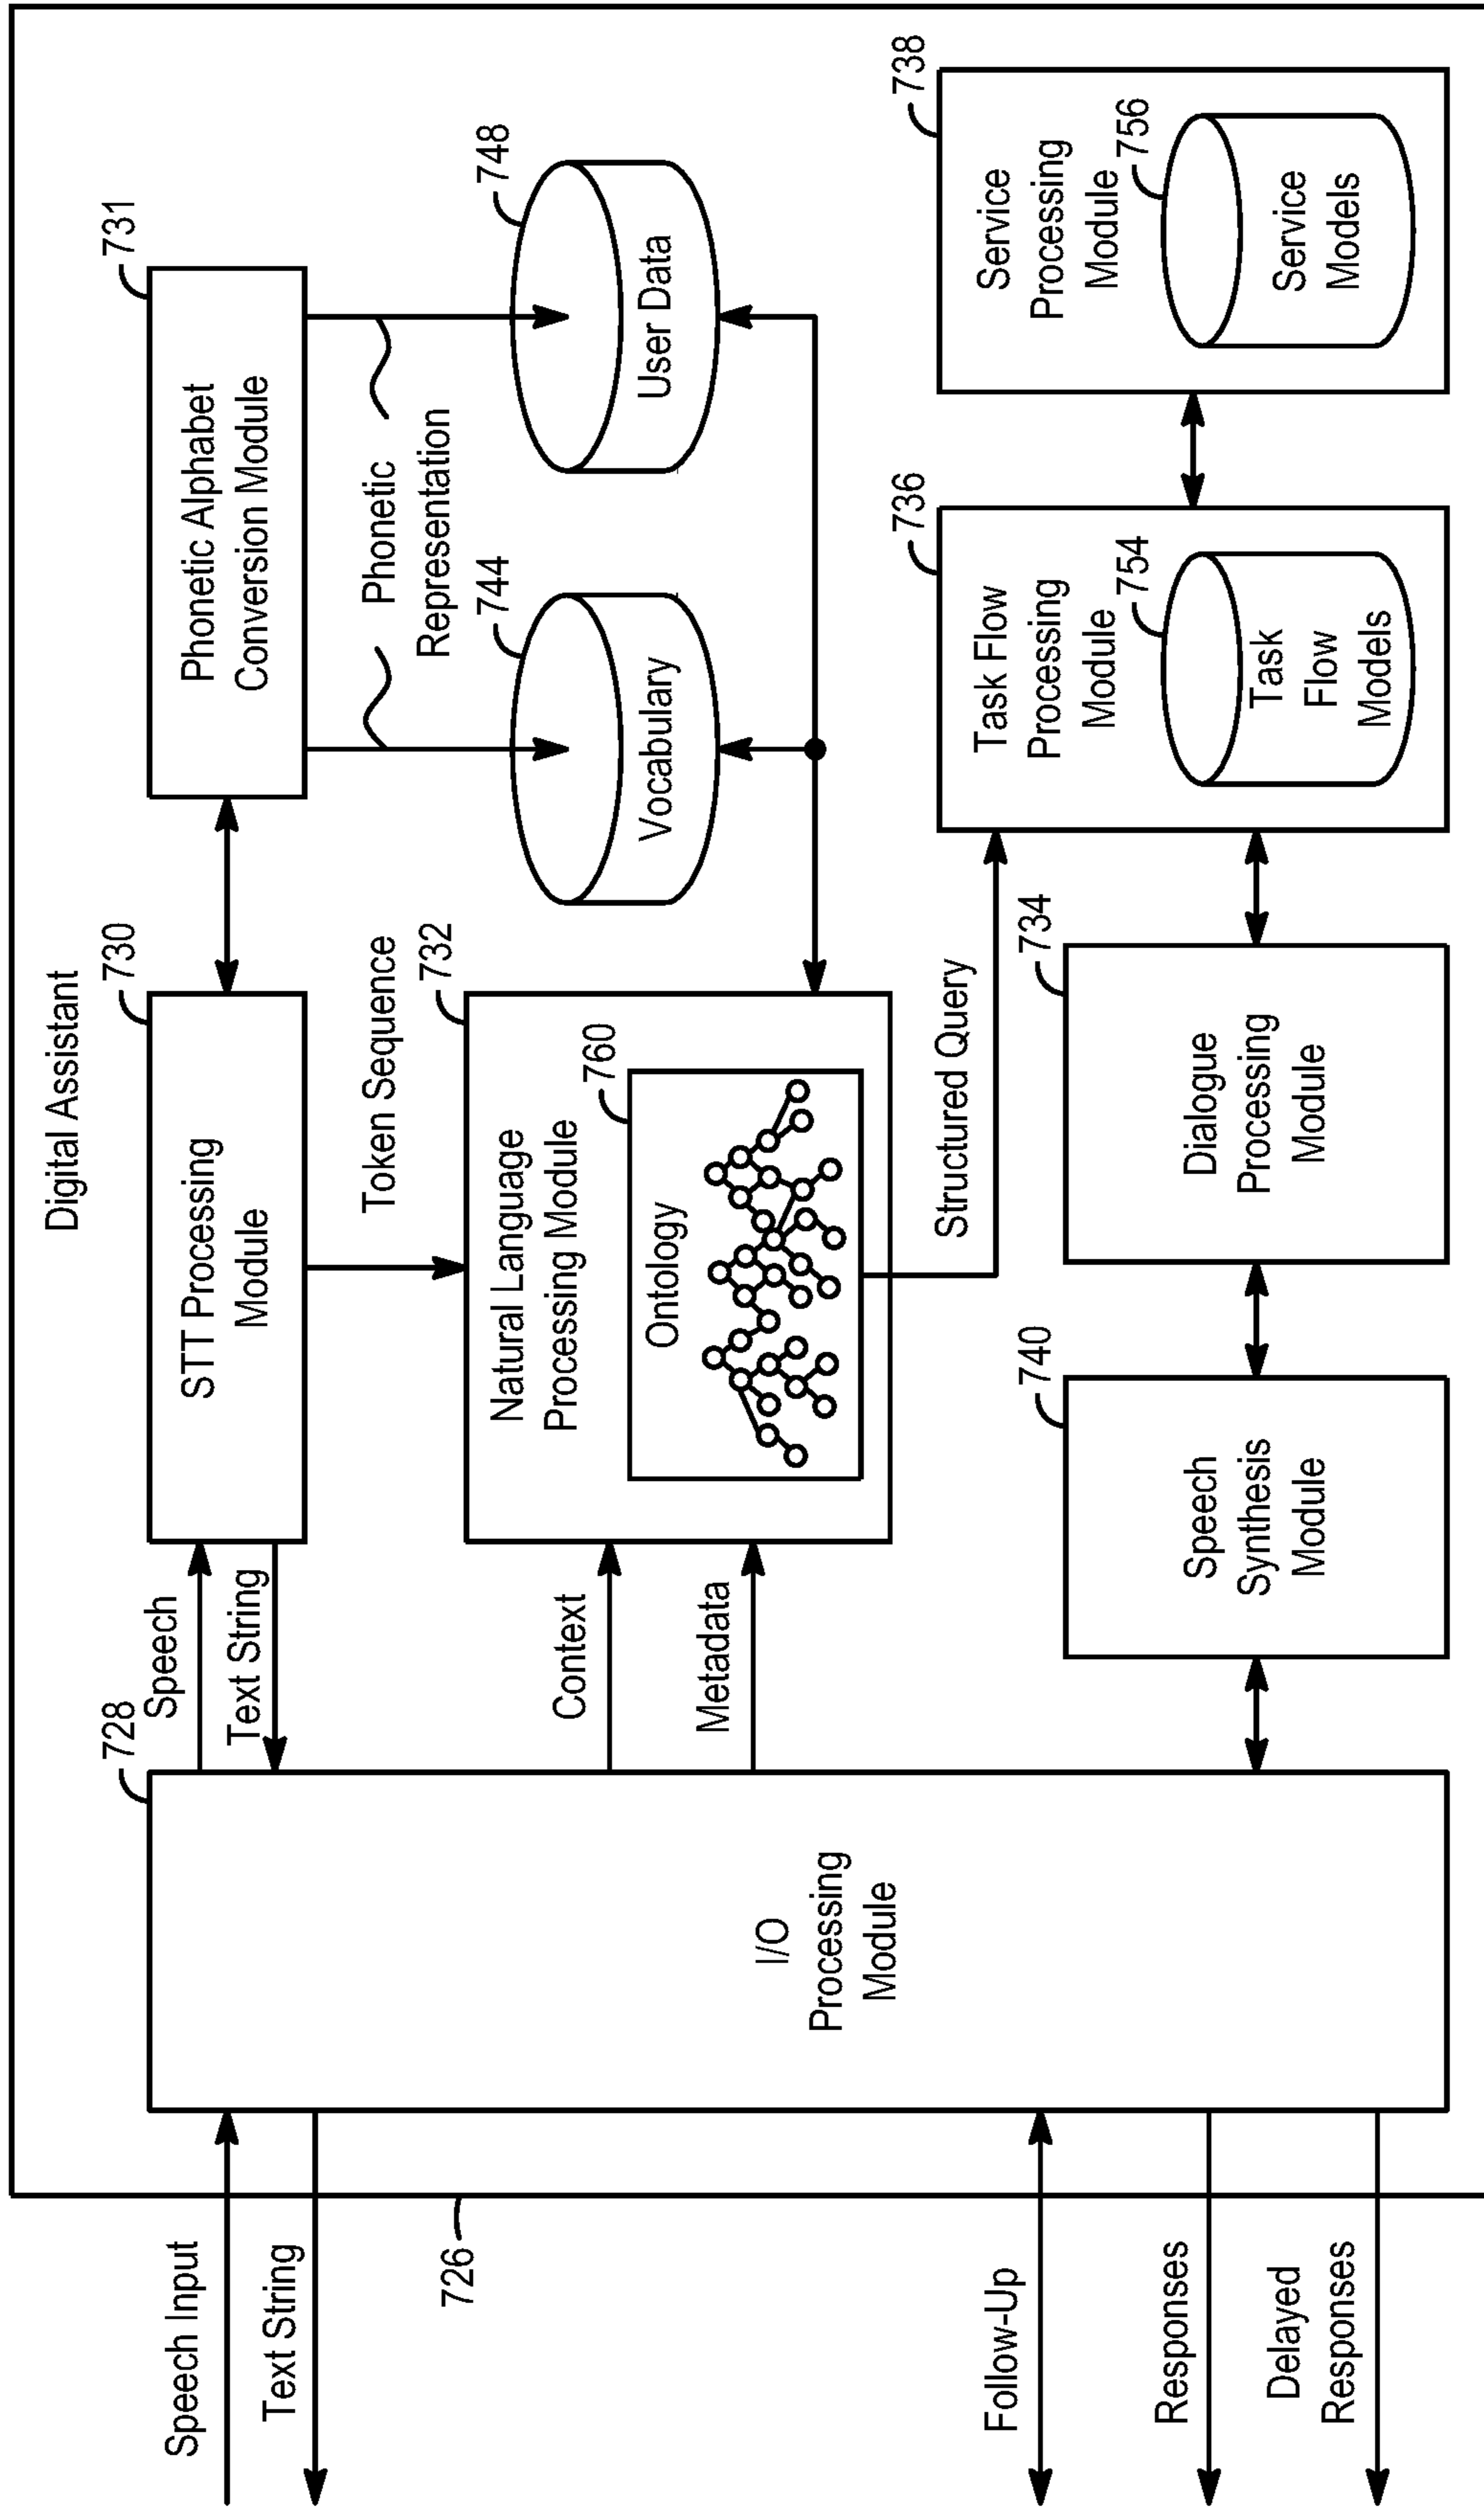


FIG. 7B

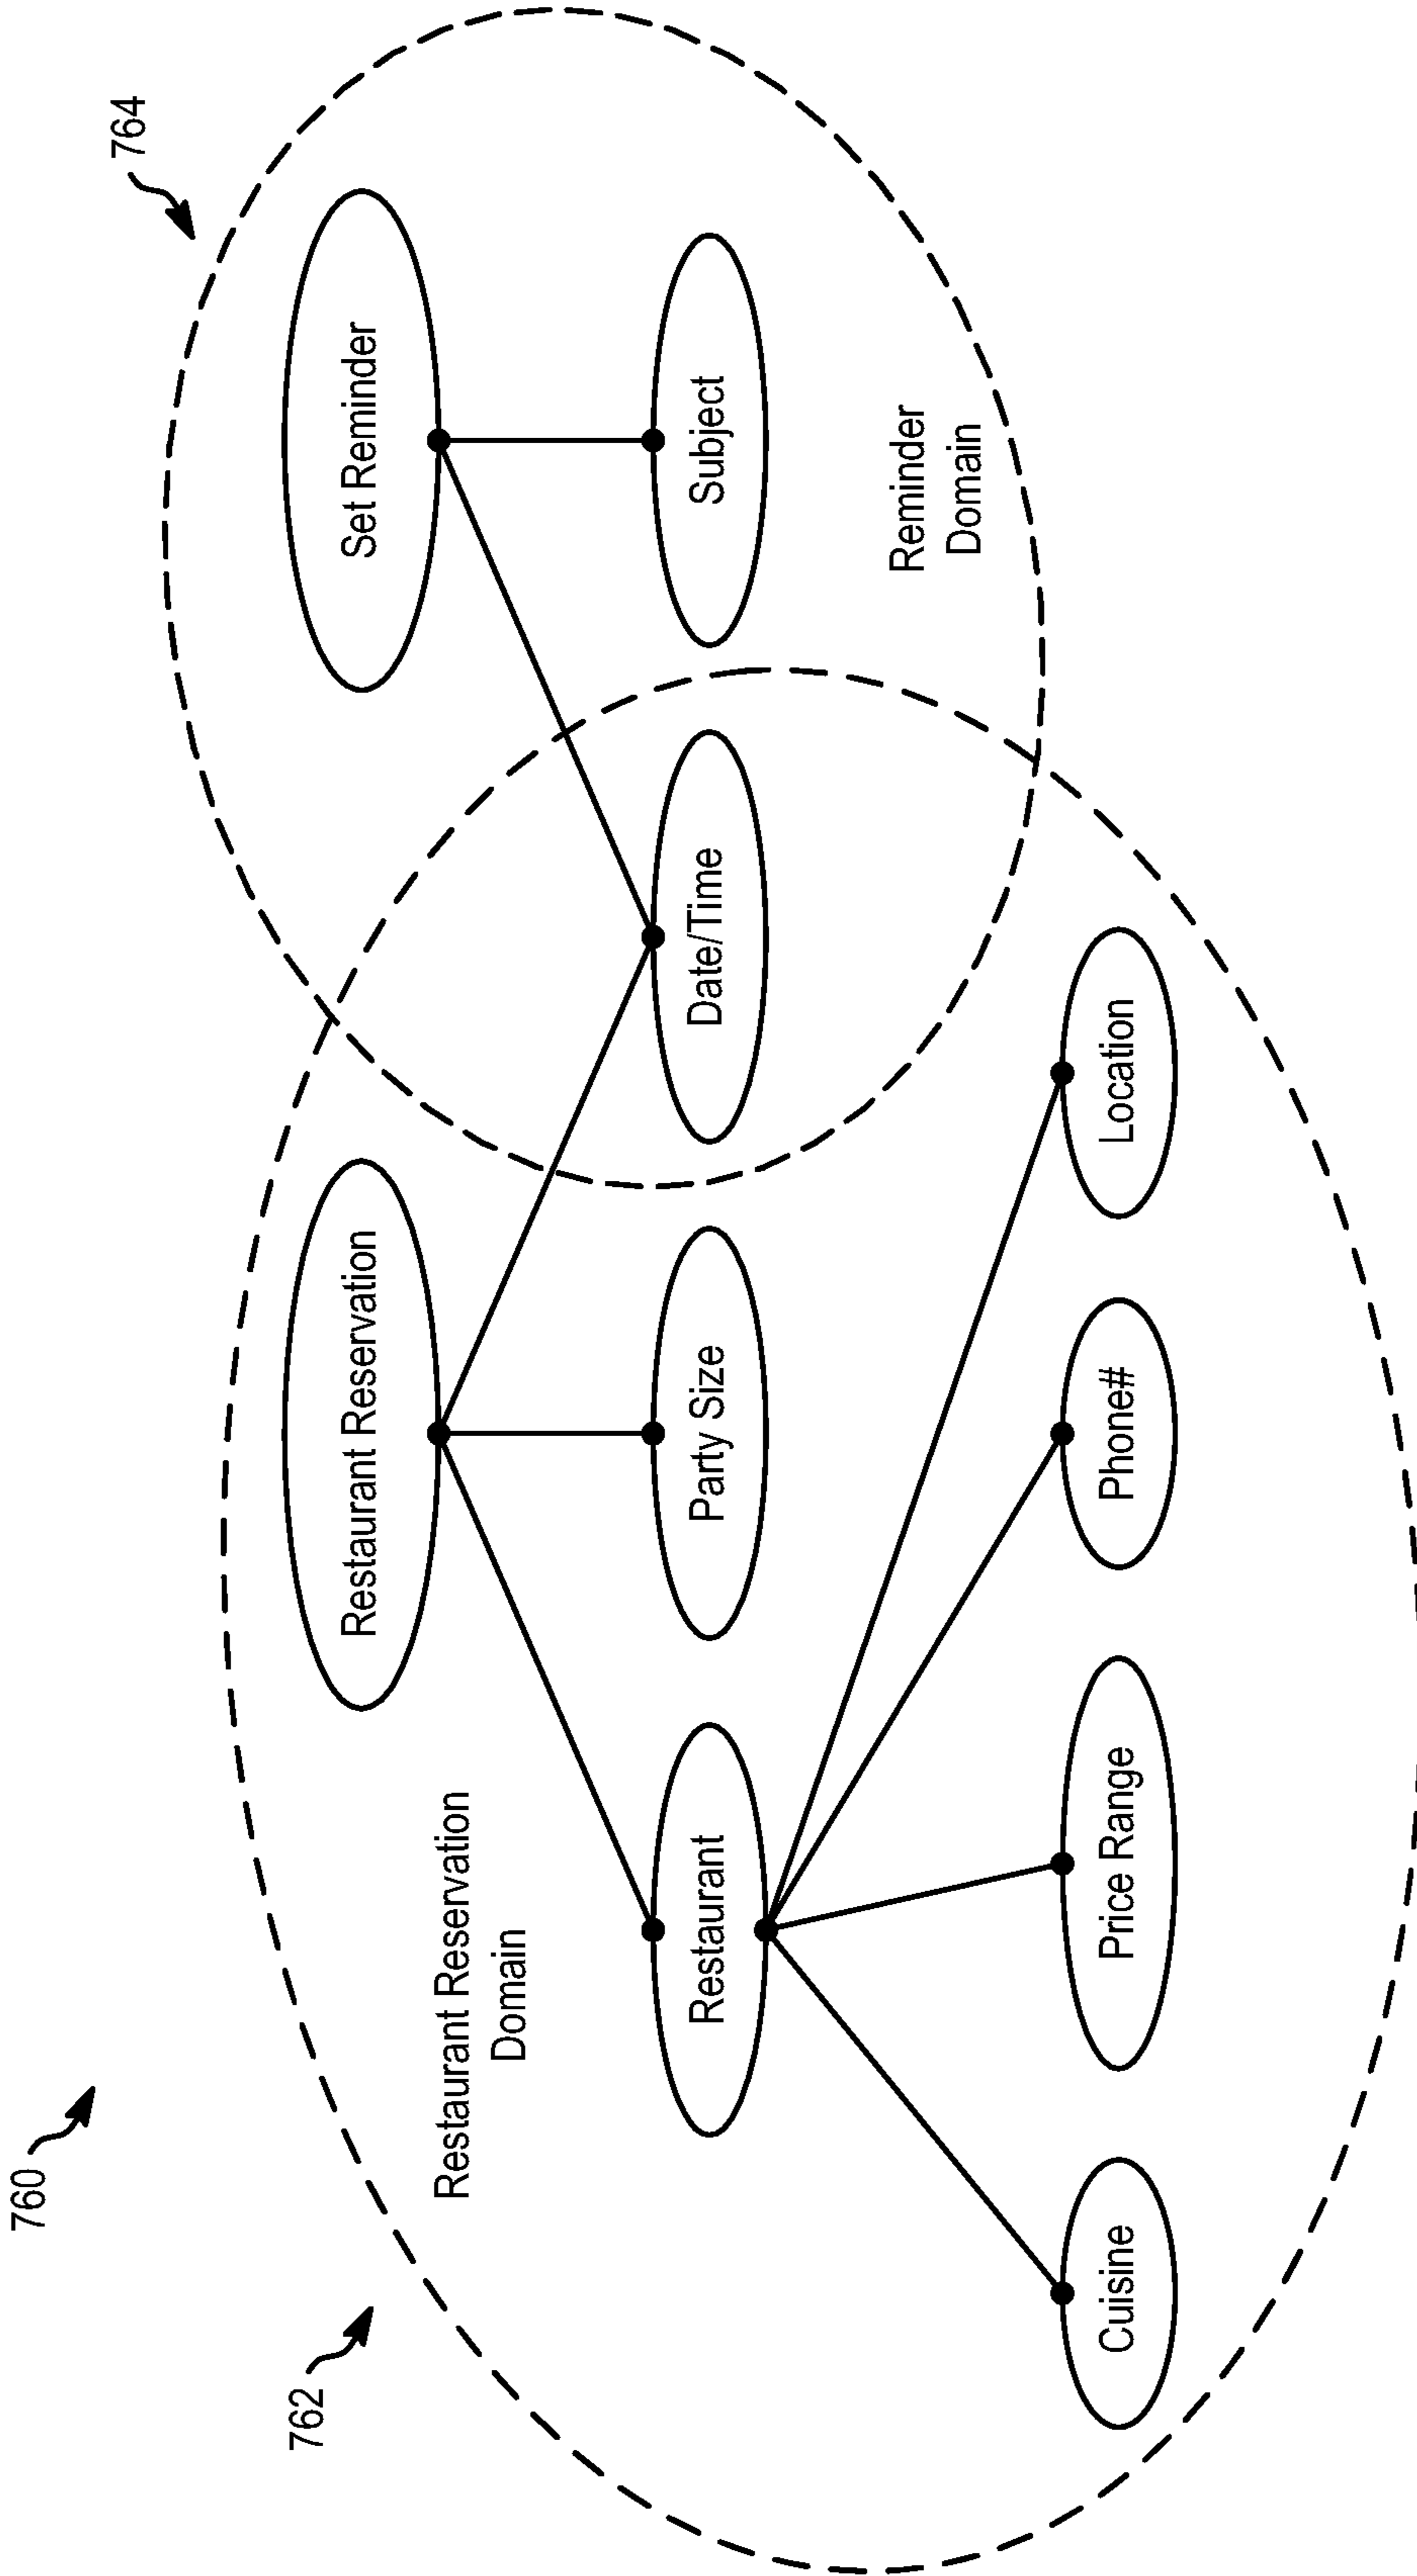


FIG. 7C




FIG. 8




FIG. 9




FIG. 10






FIG. 11A

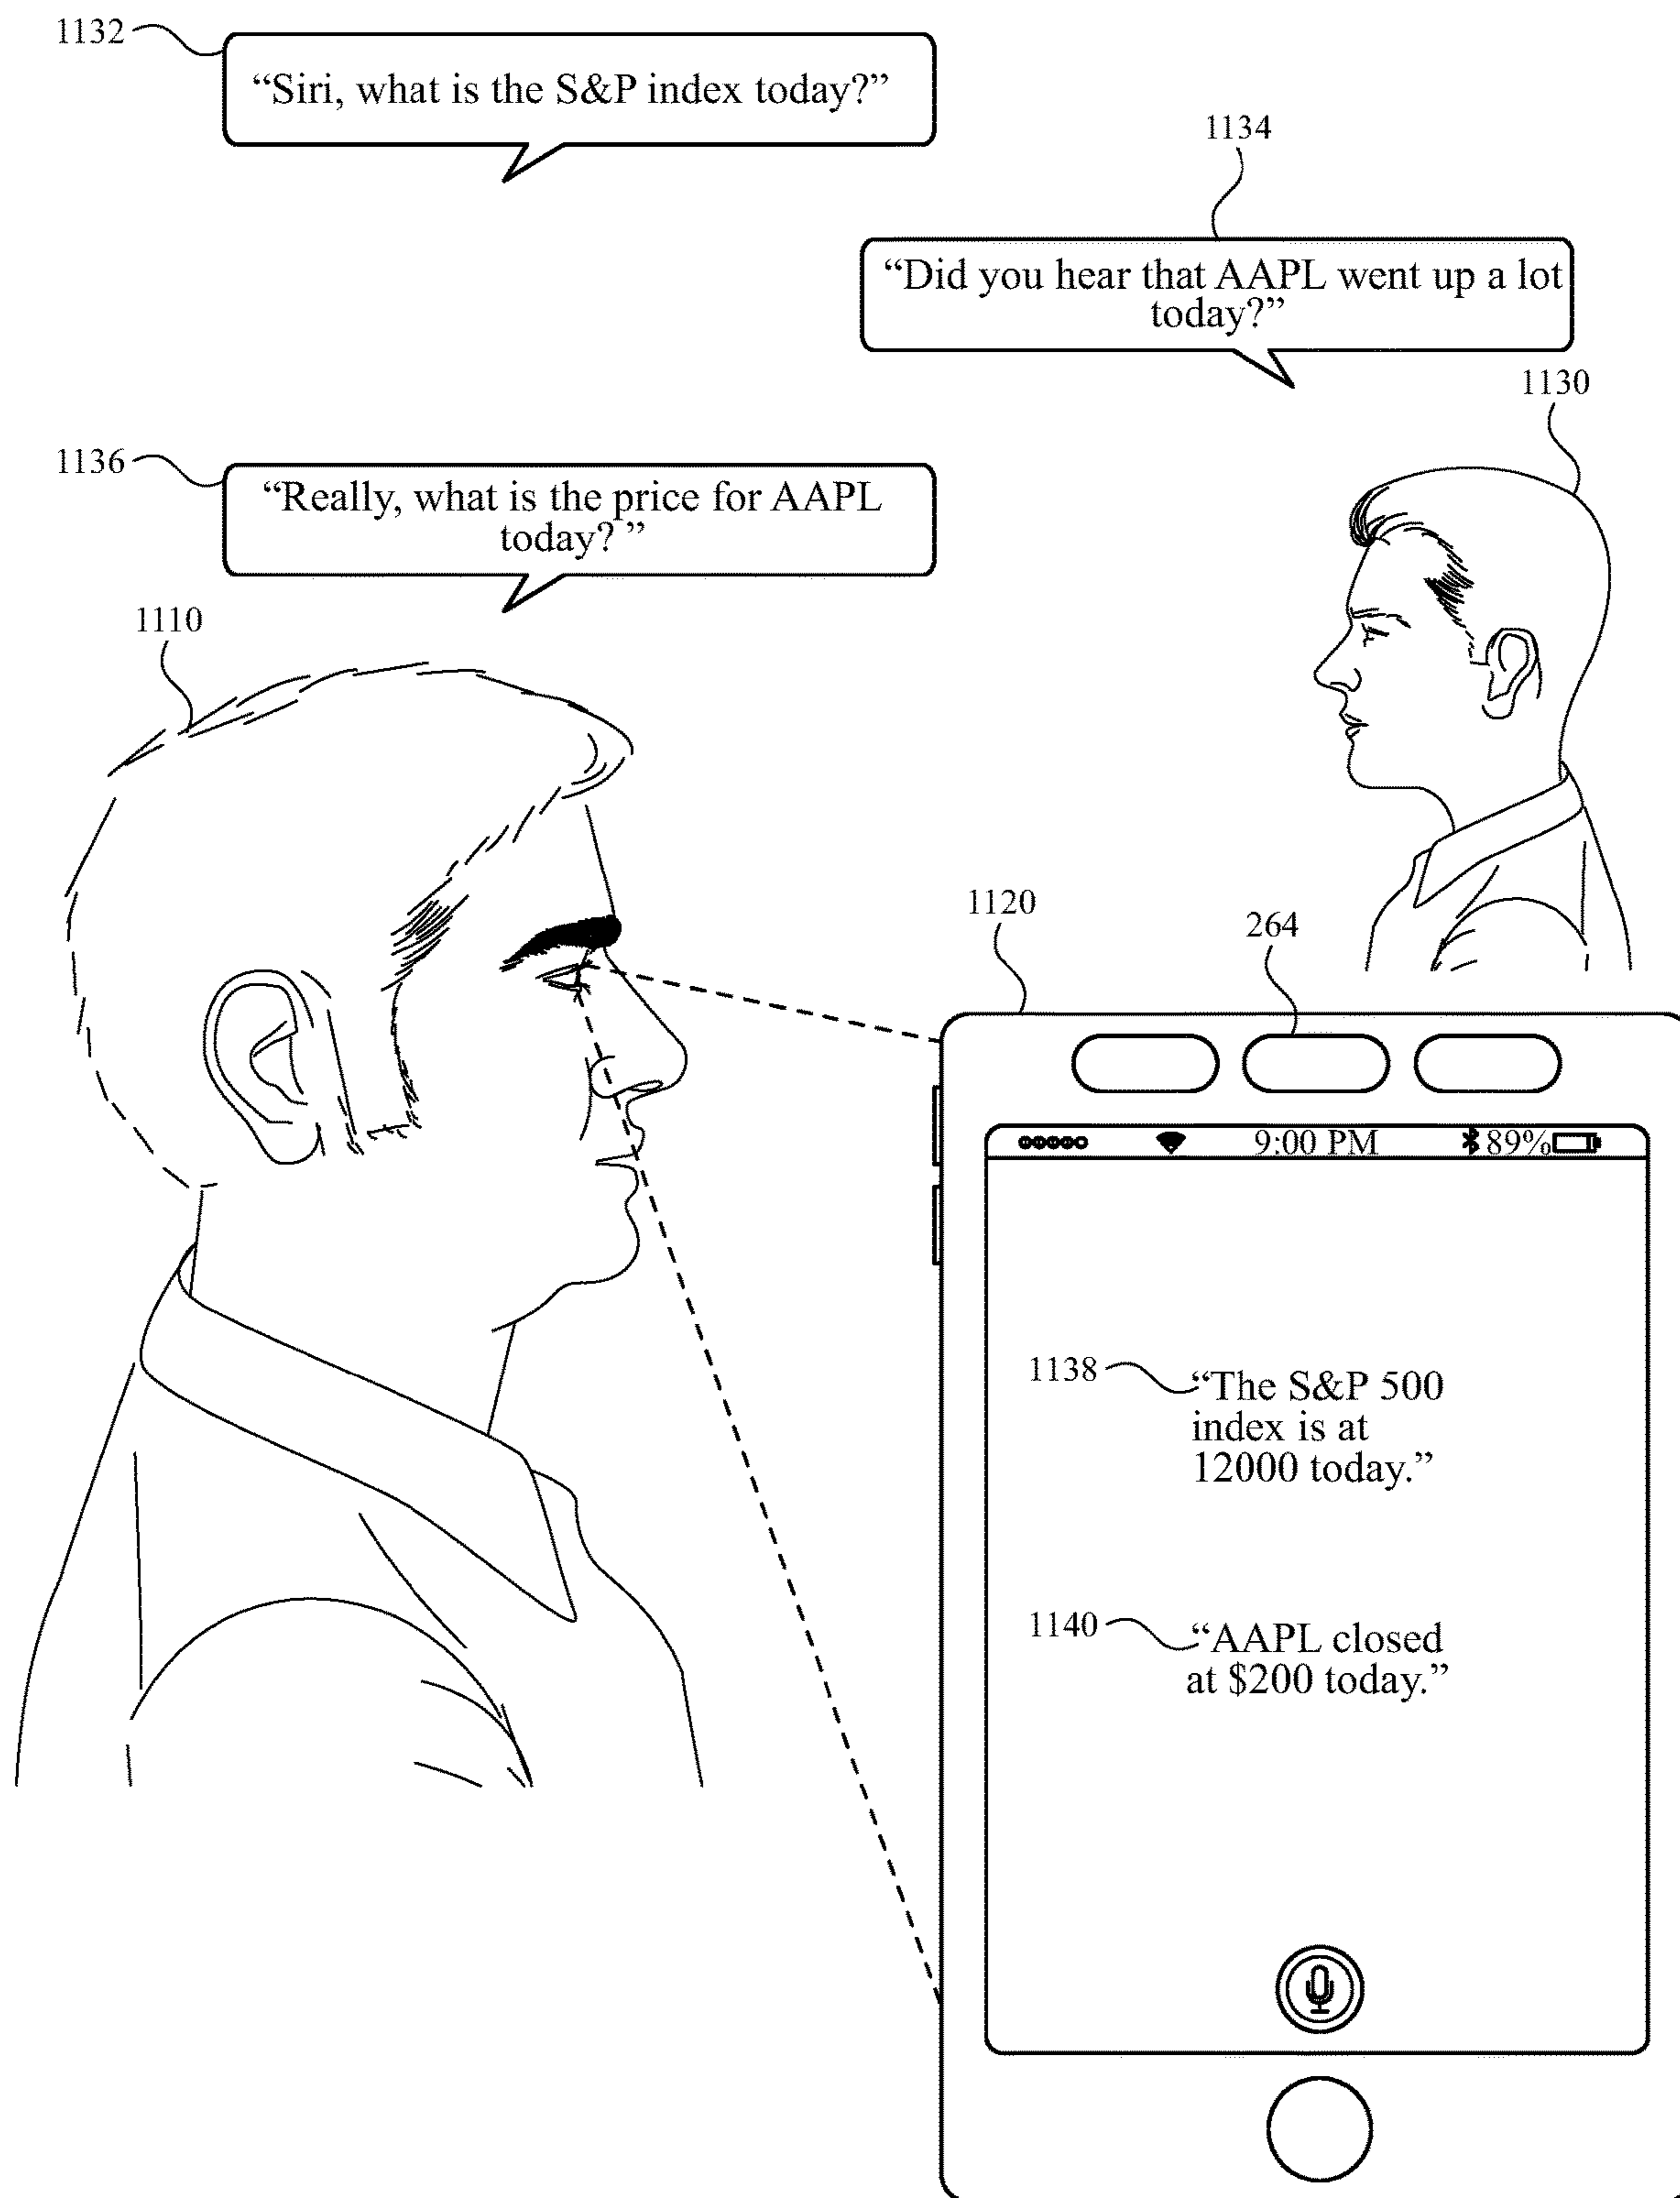


FIG. 11B




FIG. 12A




FIG. 12B




FIG. 12C




FIG. 12D




FIG. 13A




FIG. 13B






FIG. 14A




FIG. 14B




FIG. 14C




FIG. 14D

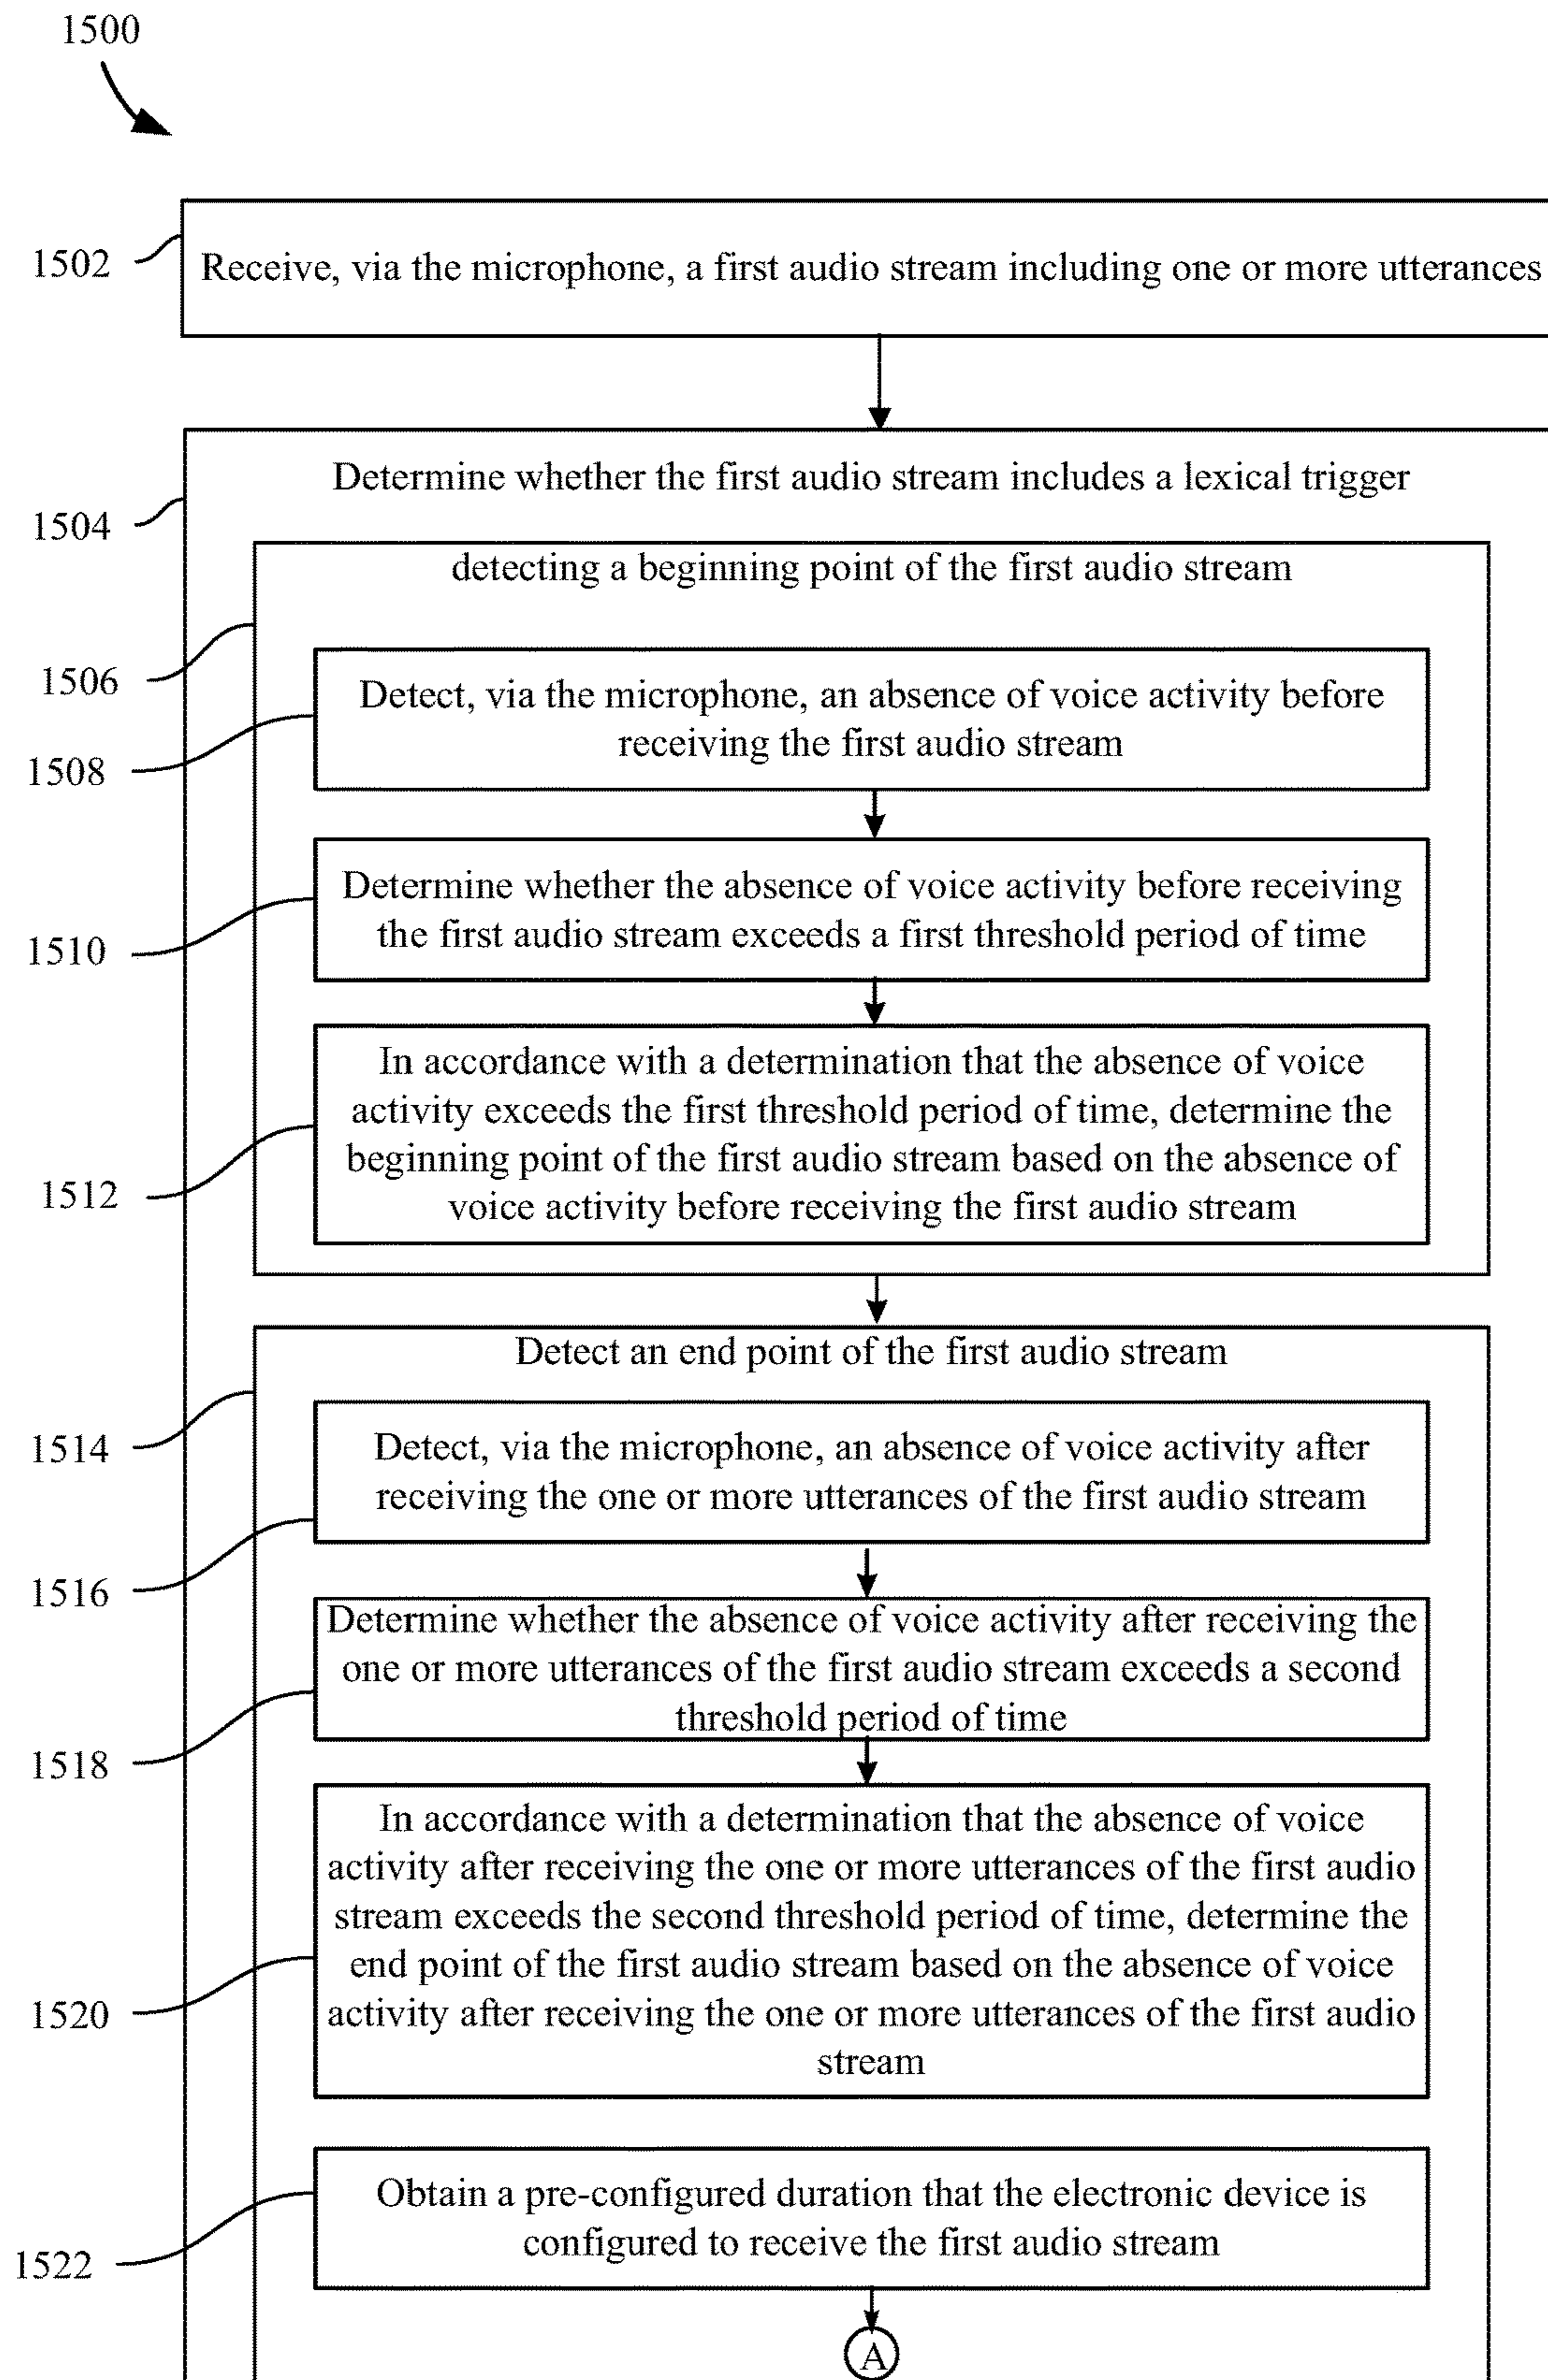


FIG. 15A

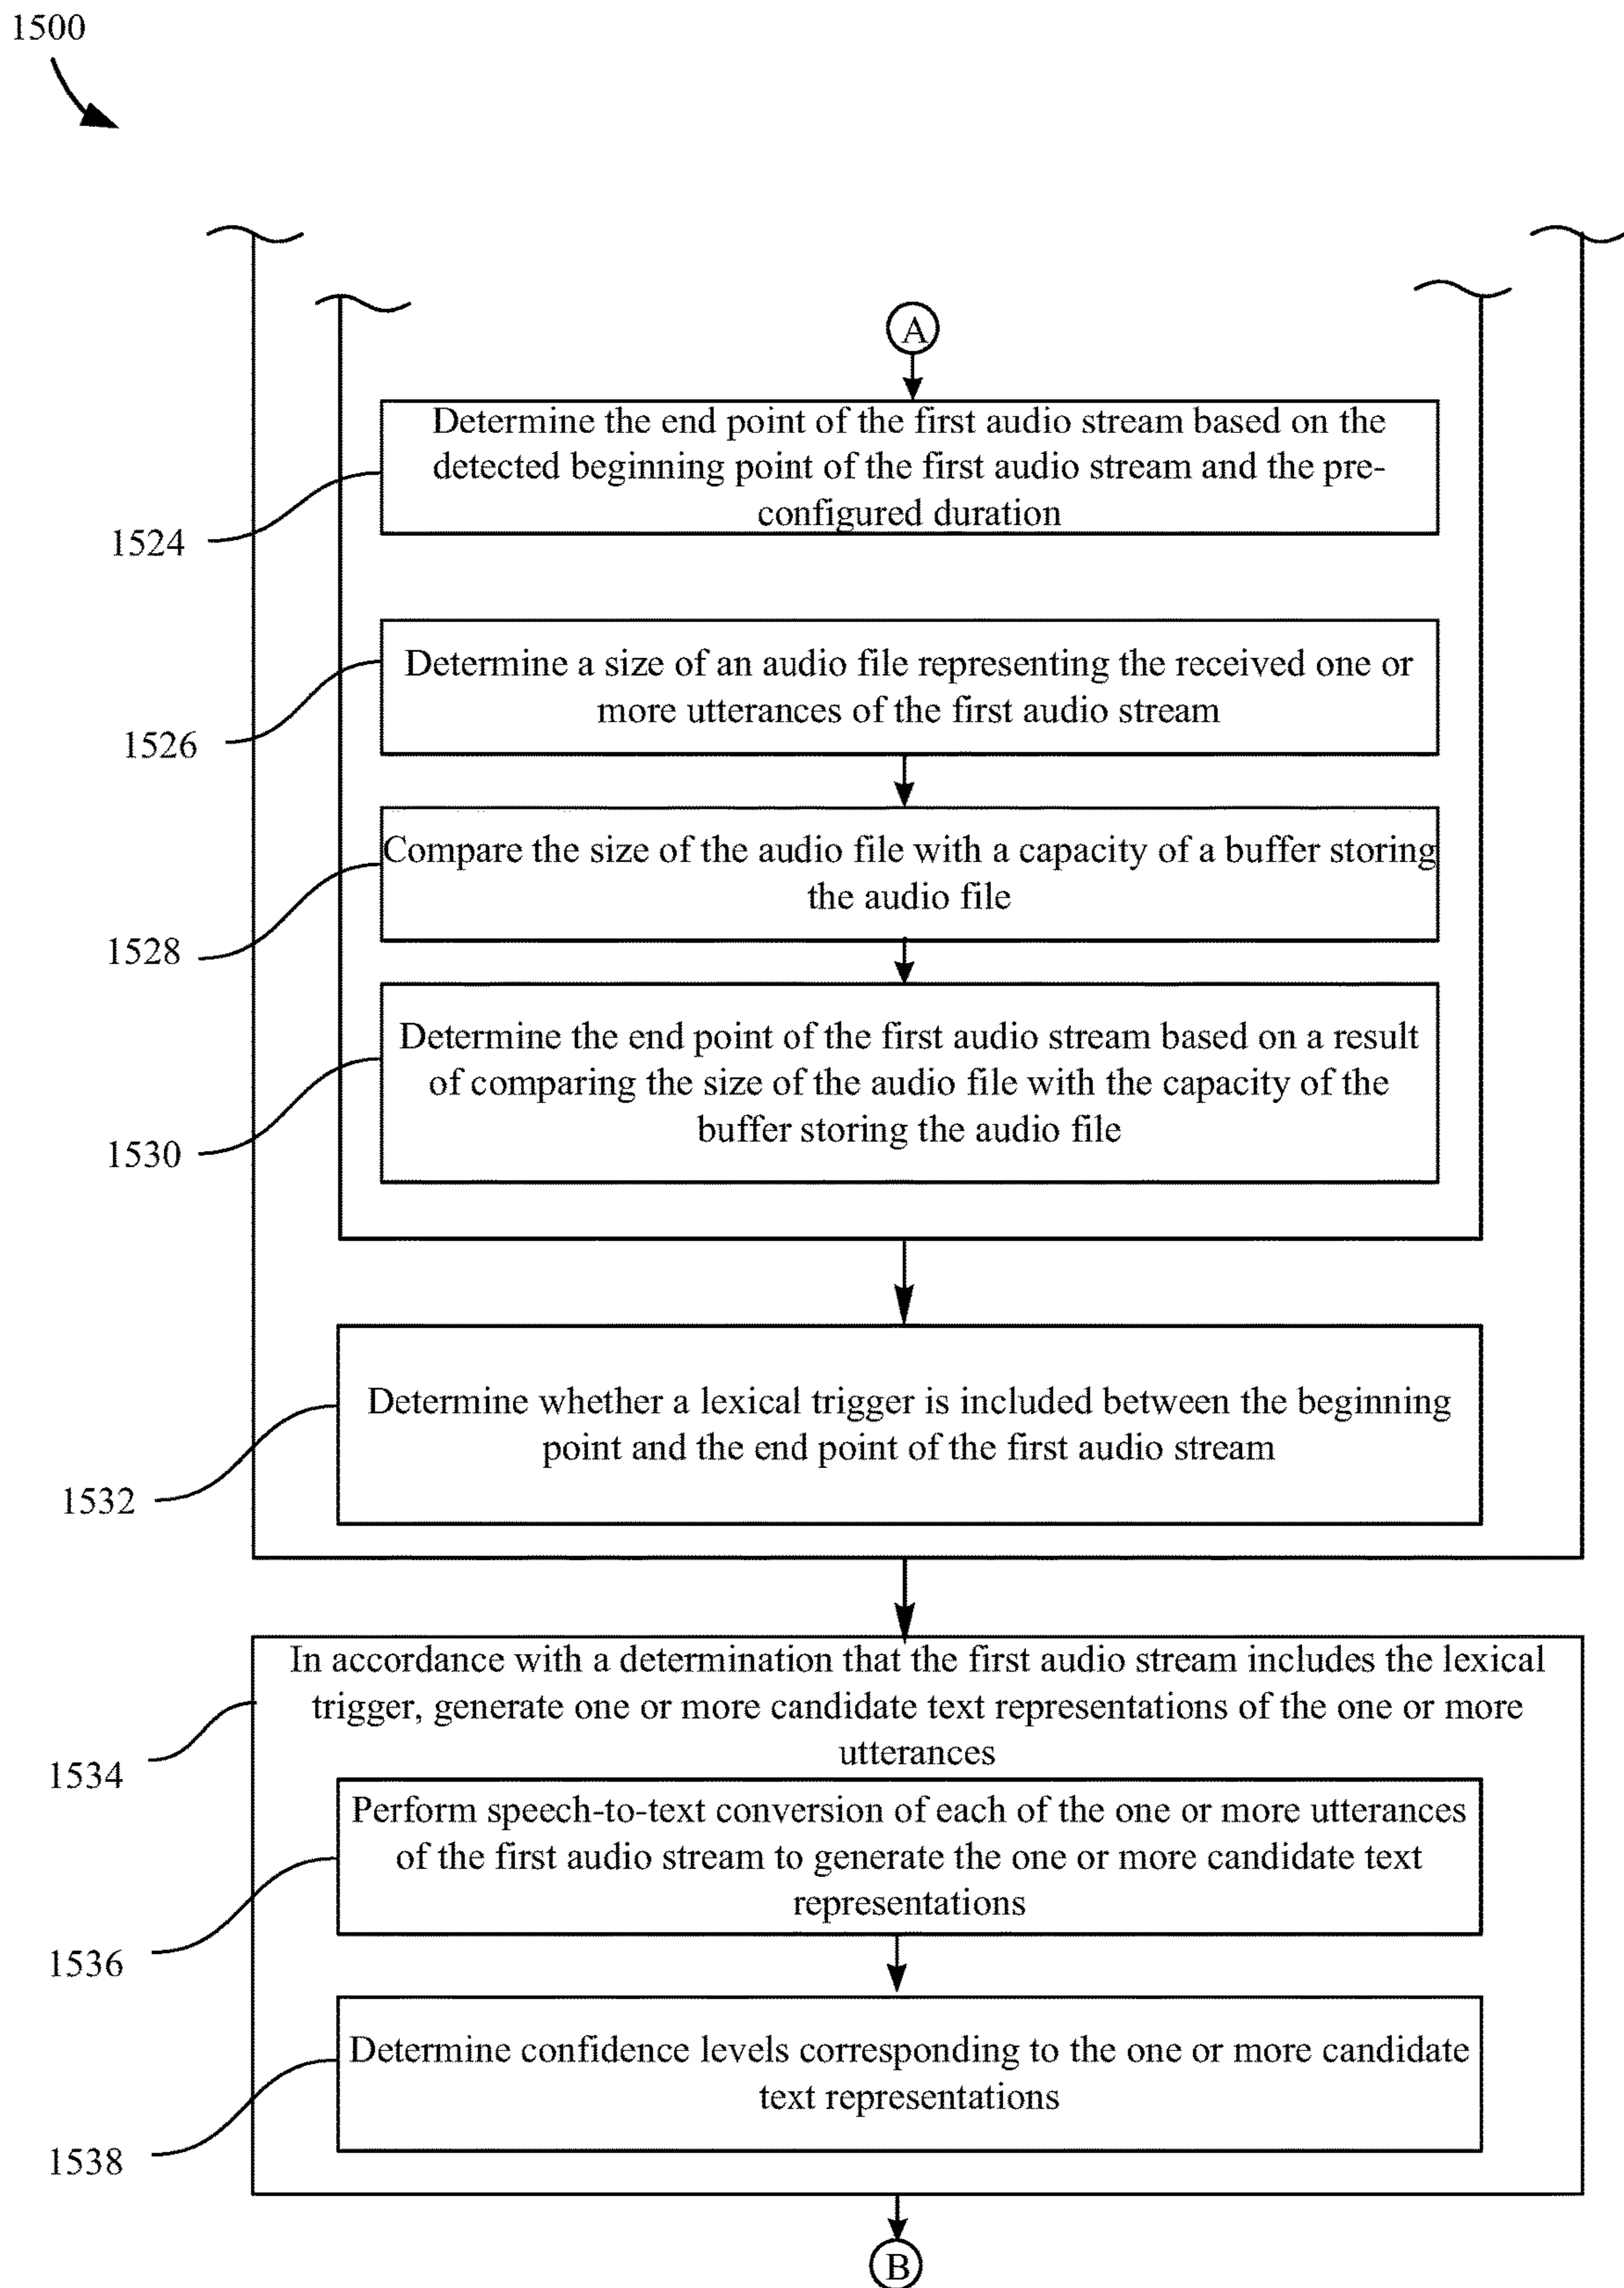


FIG. 15B

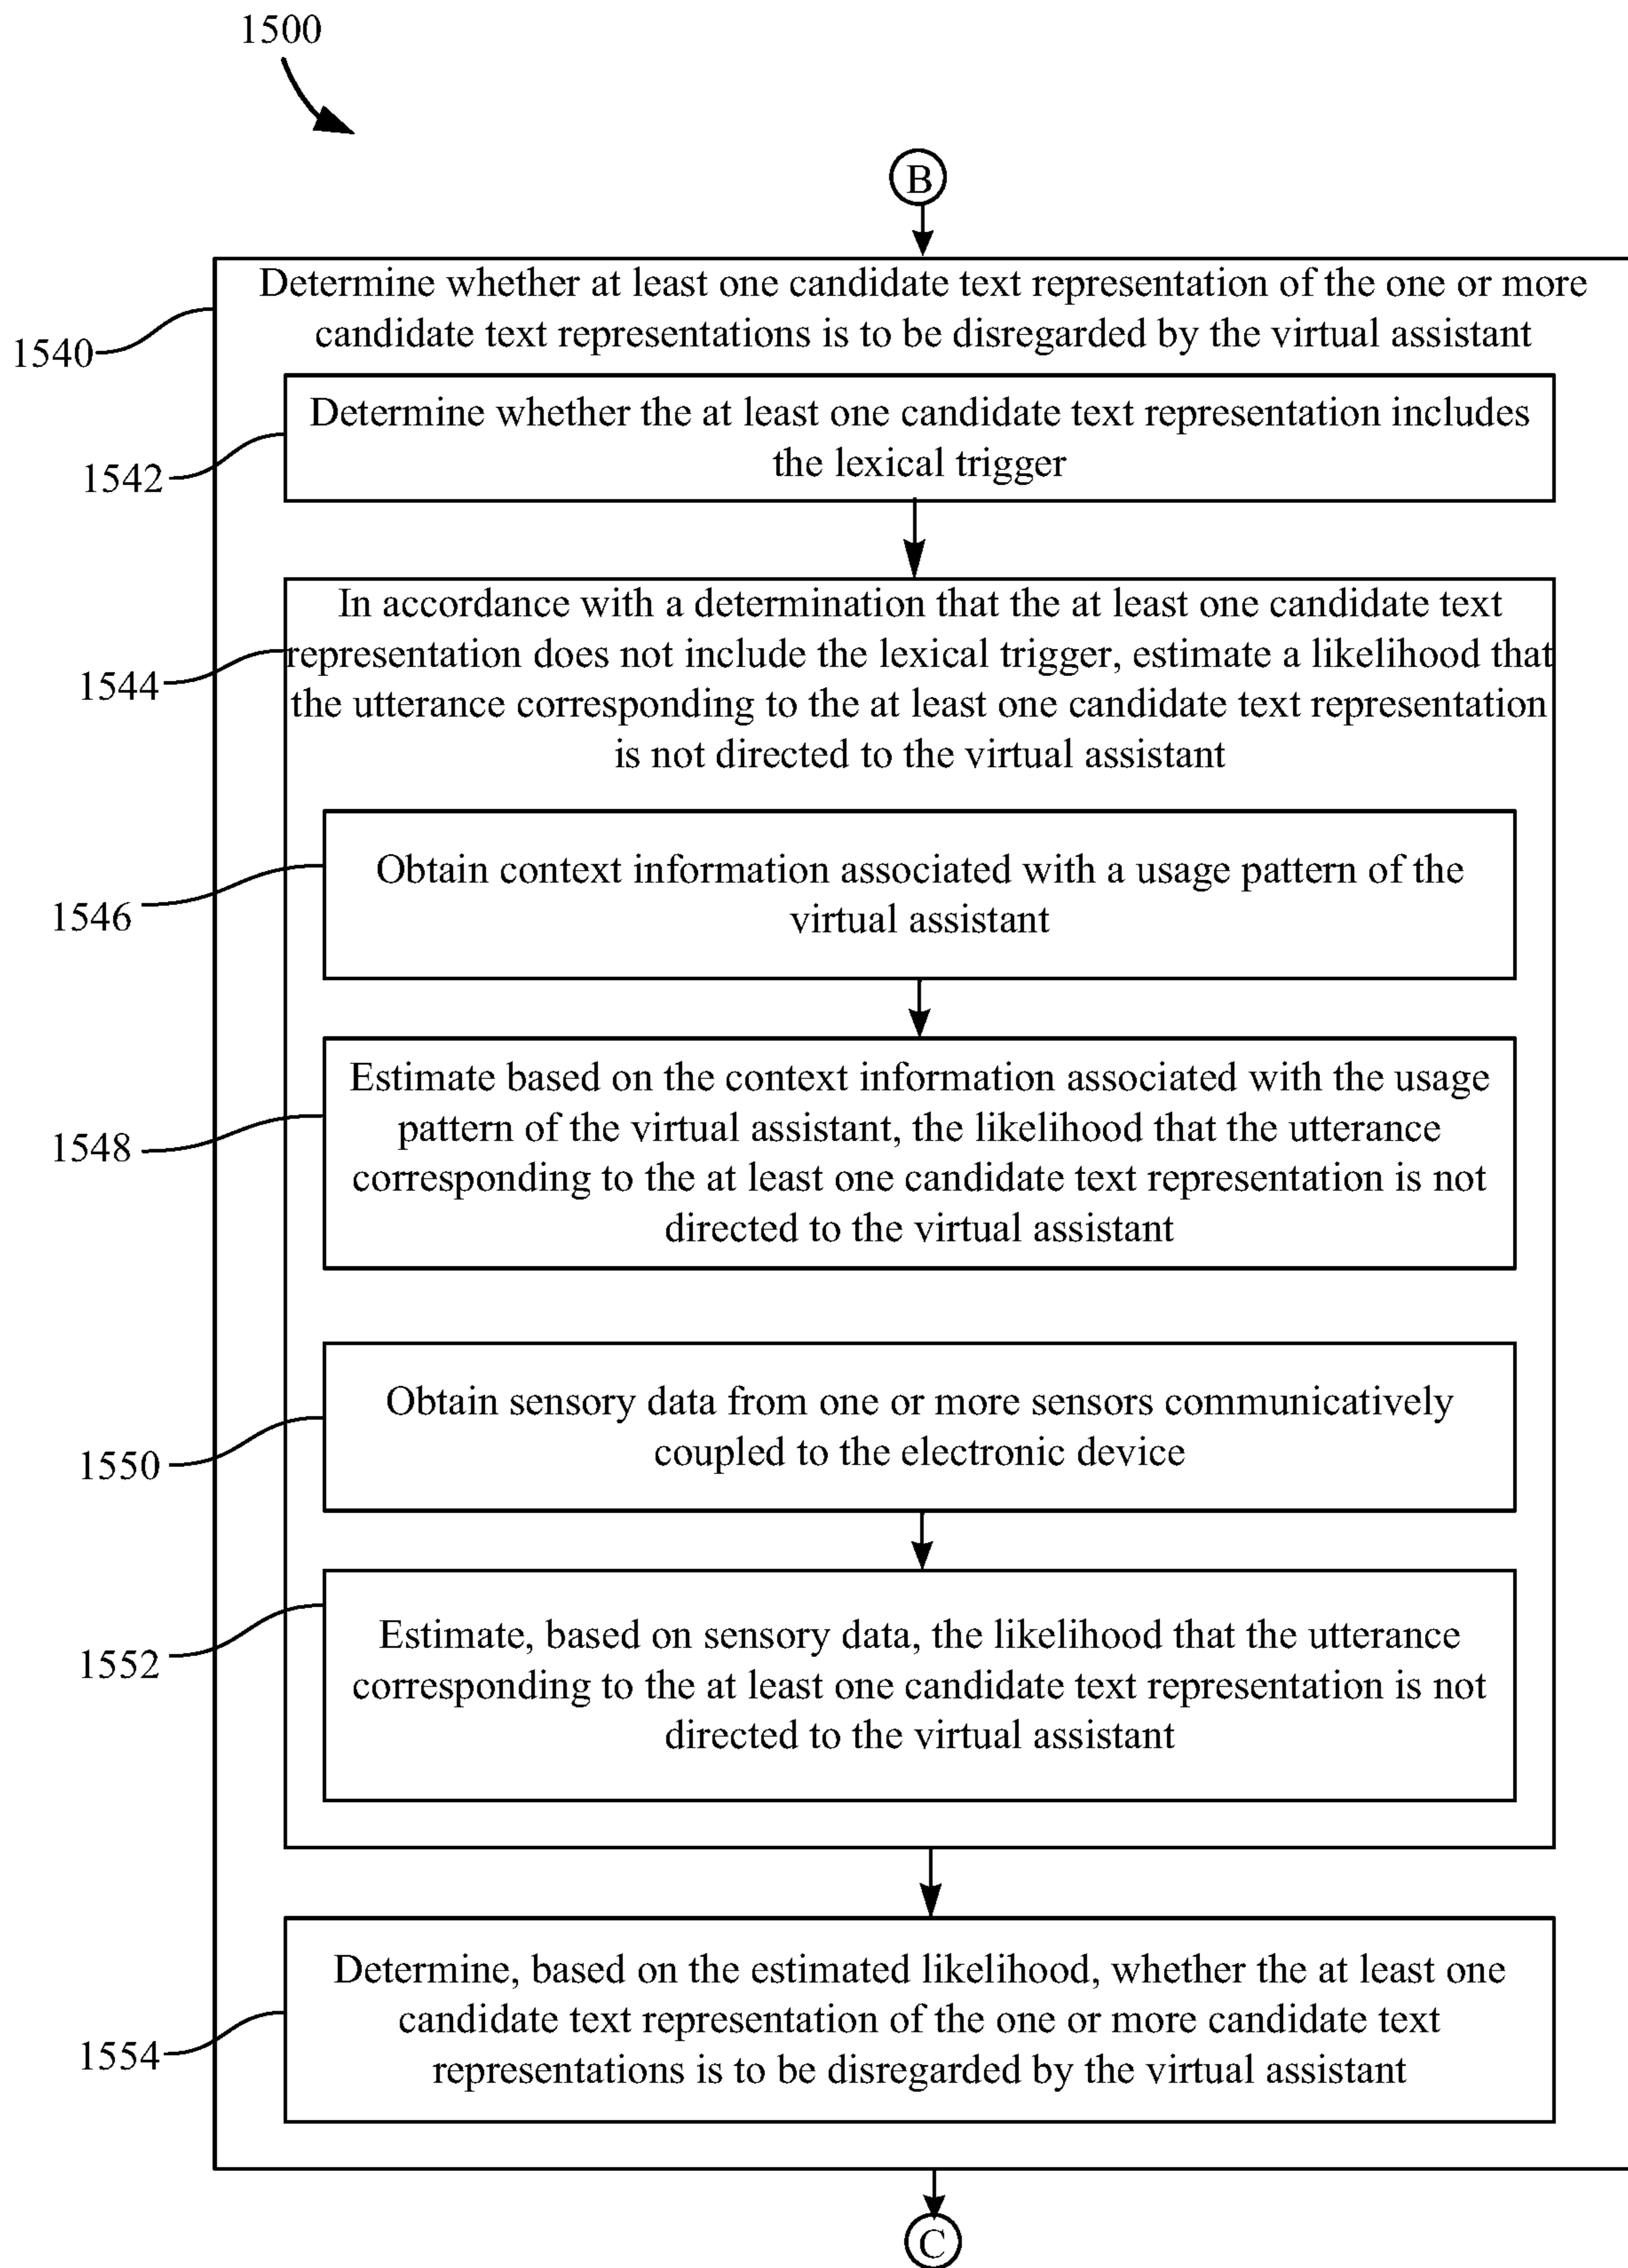


FIG. 15C

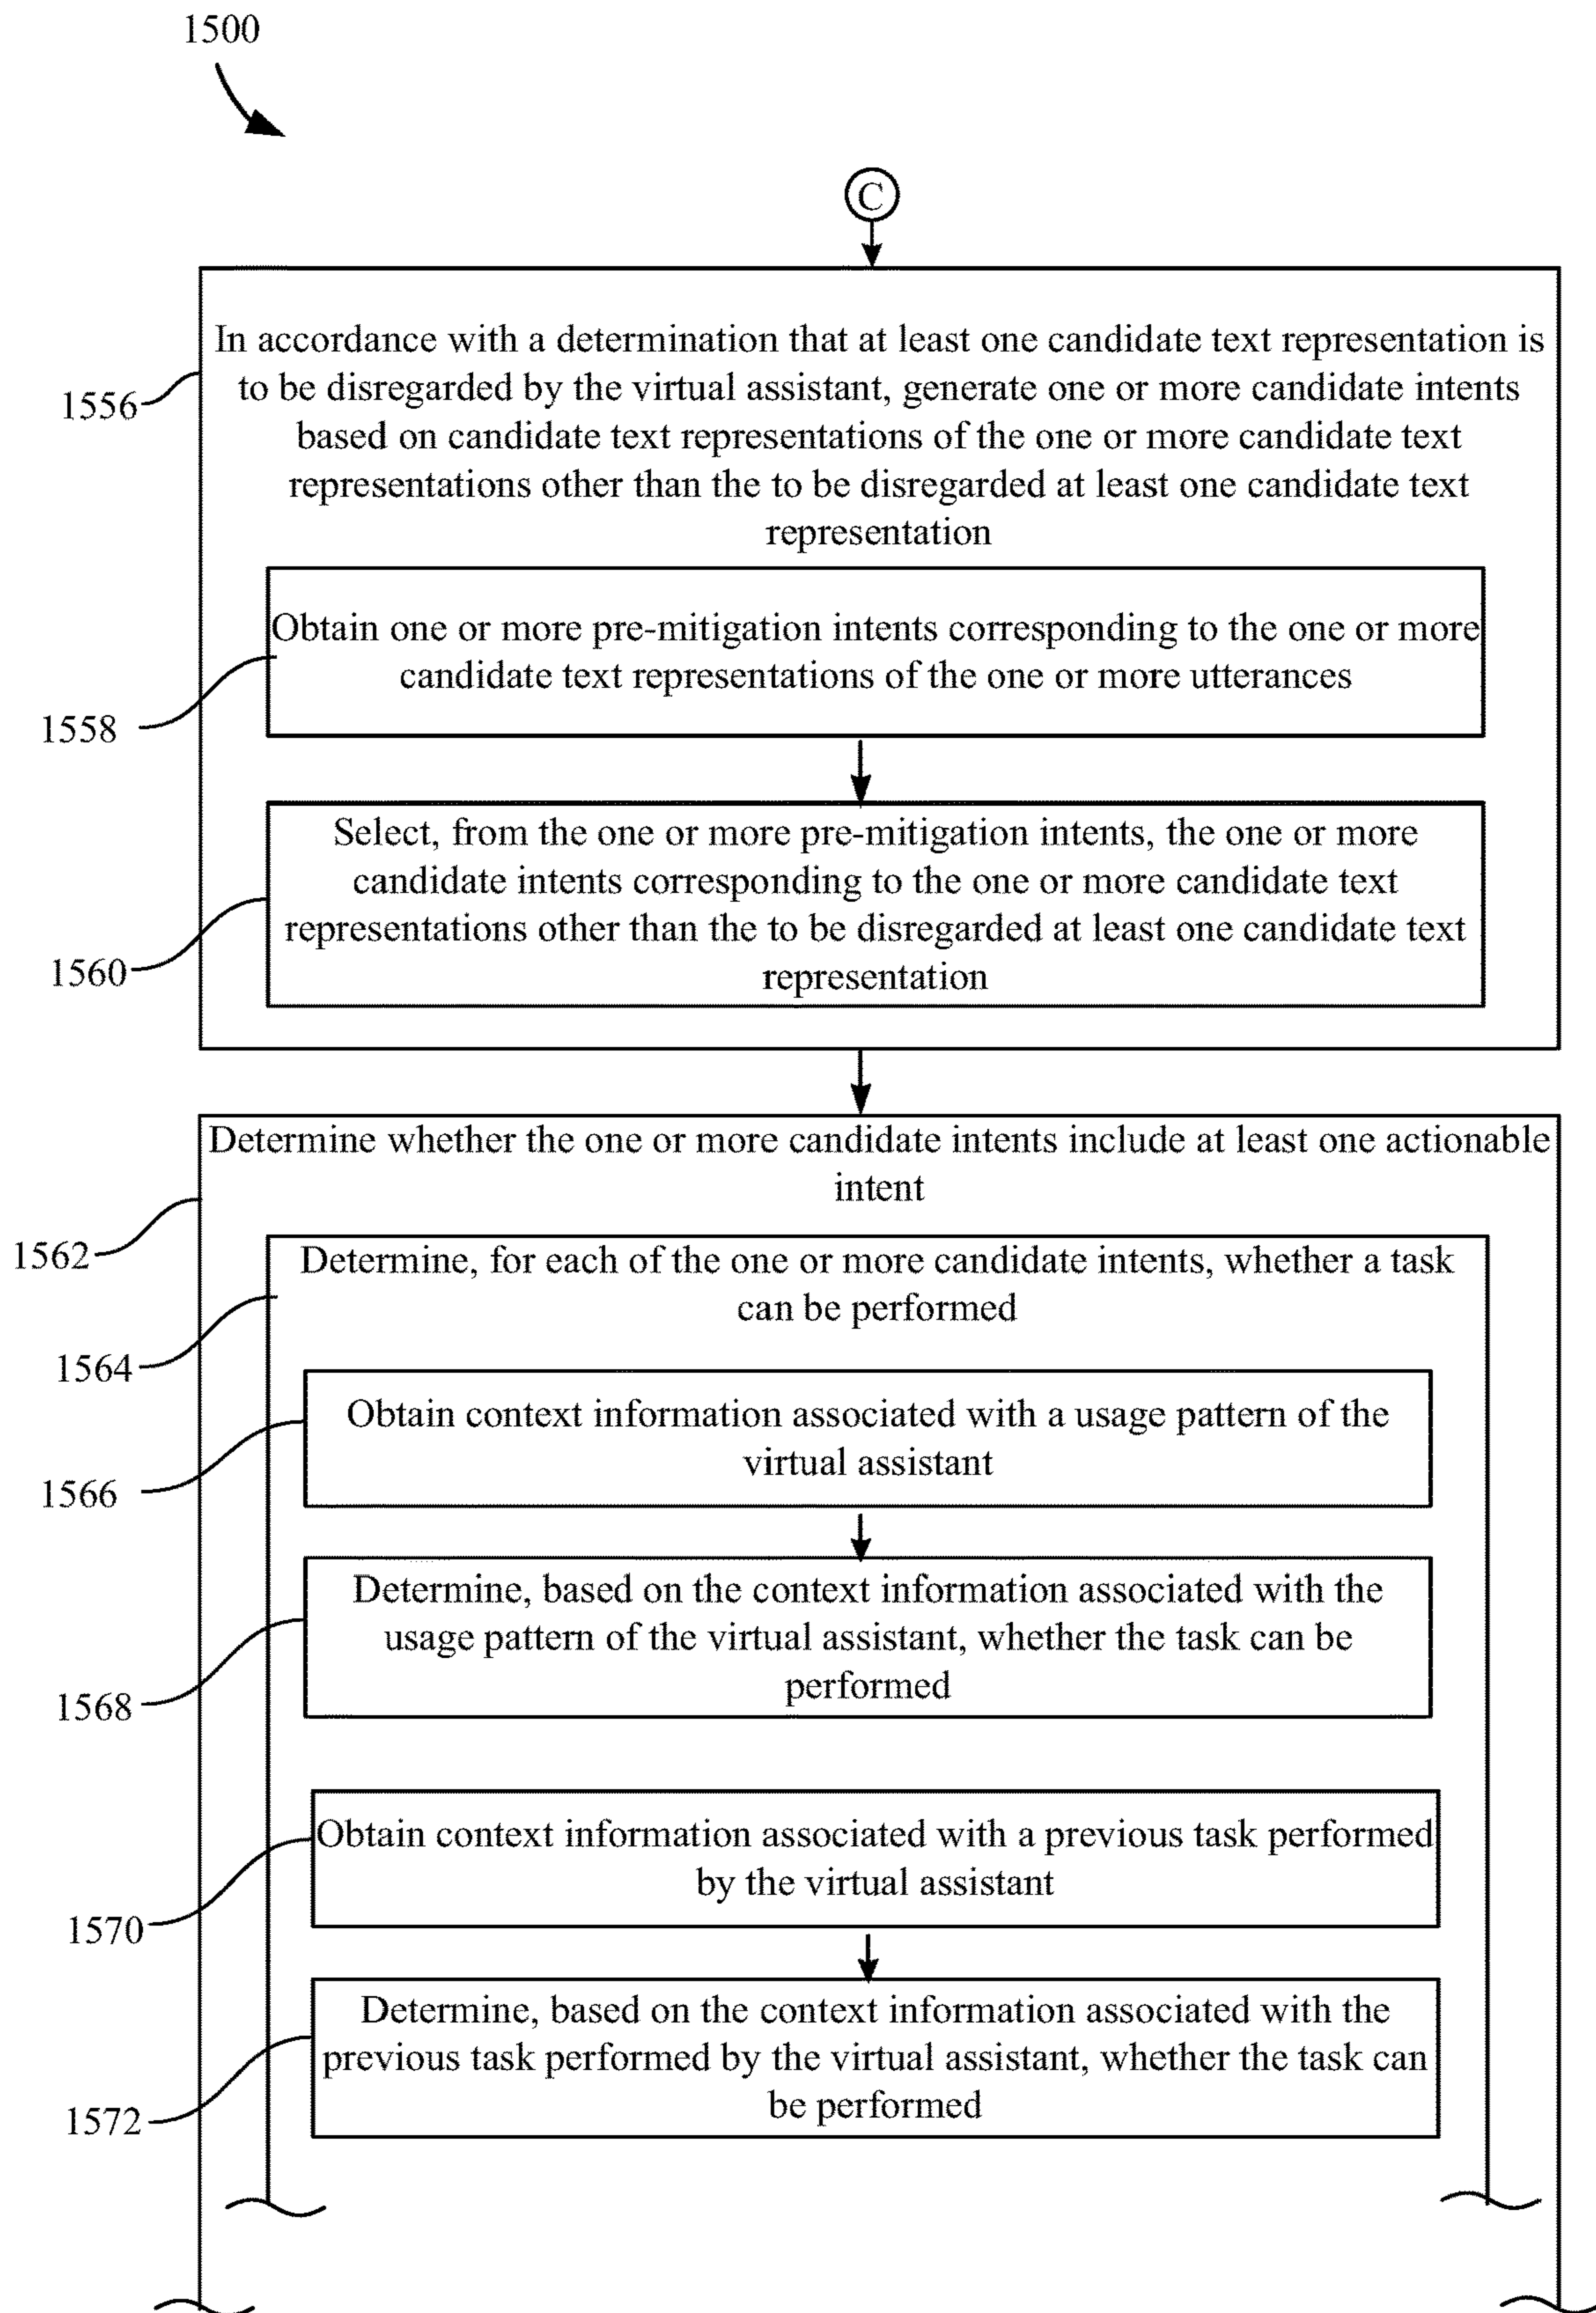


FIG. 15D



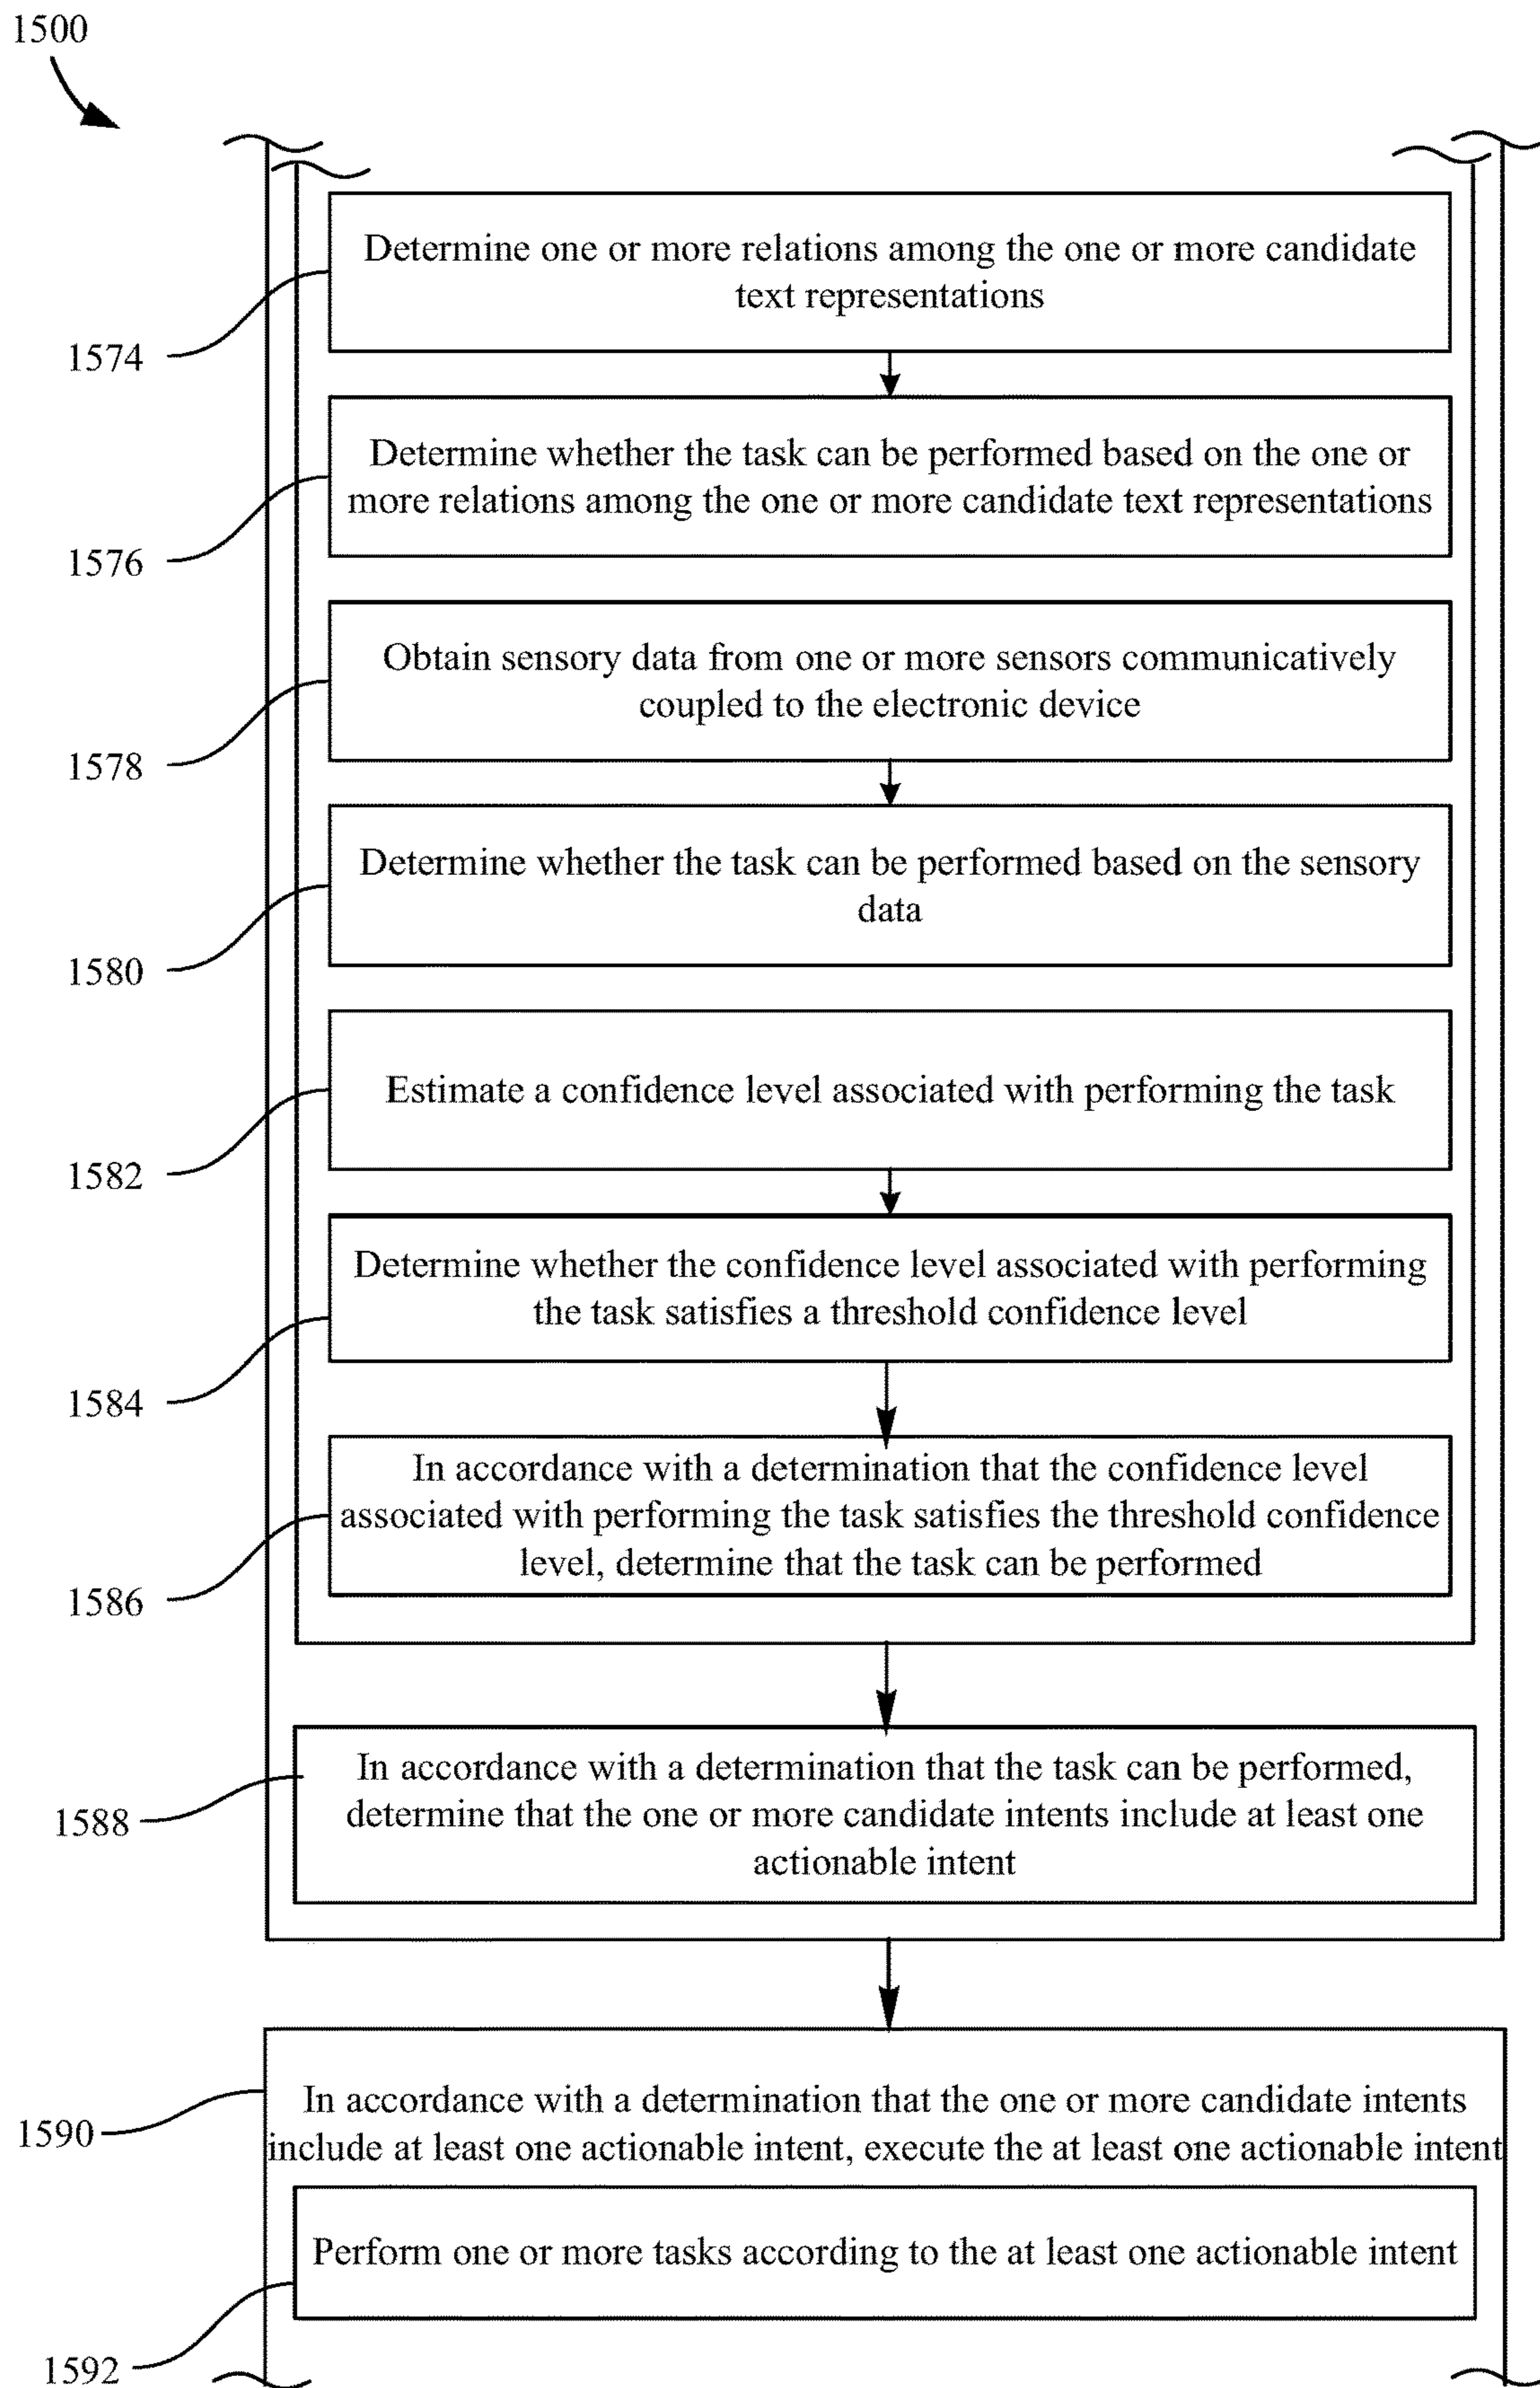


FIG. 15E

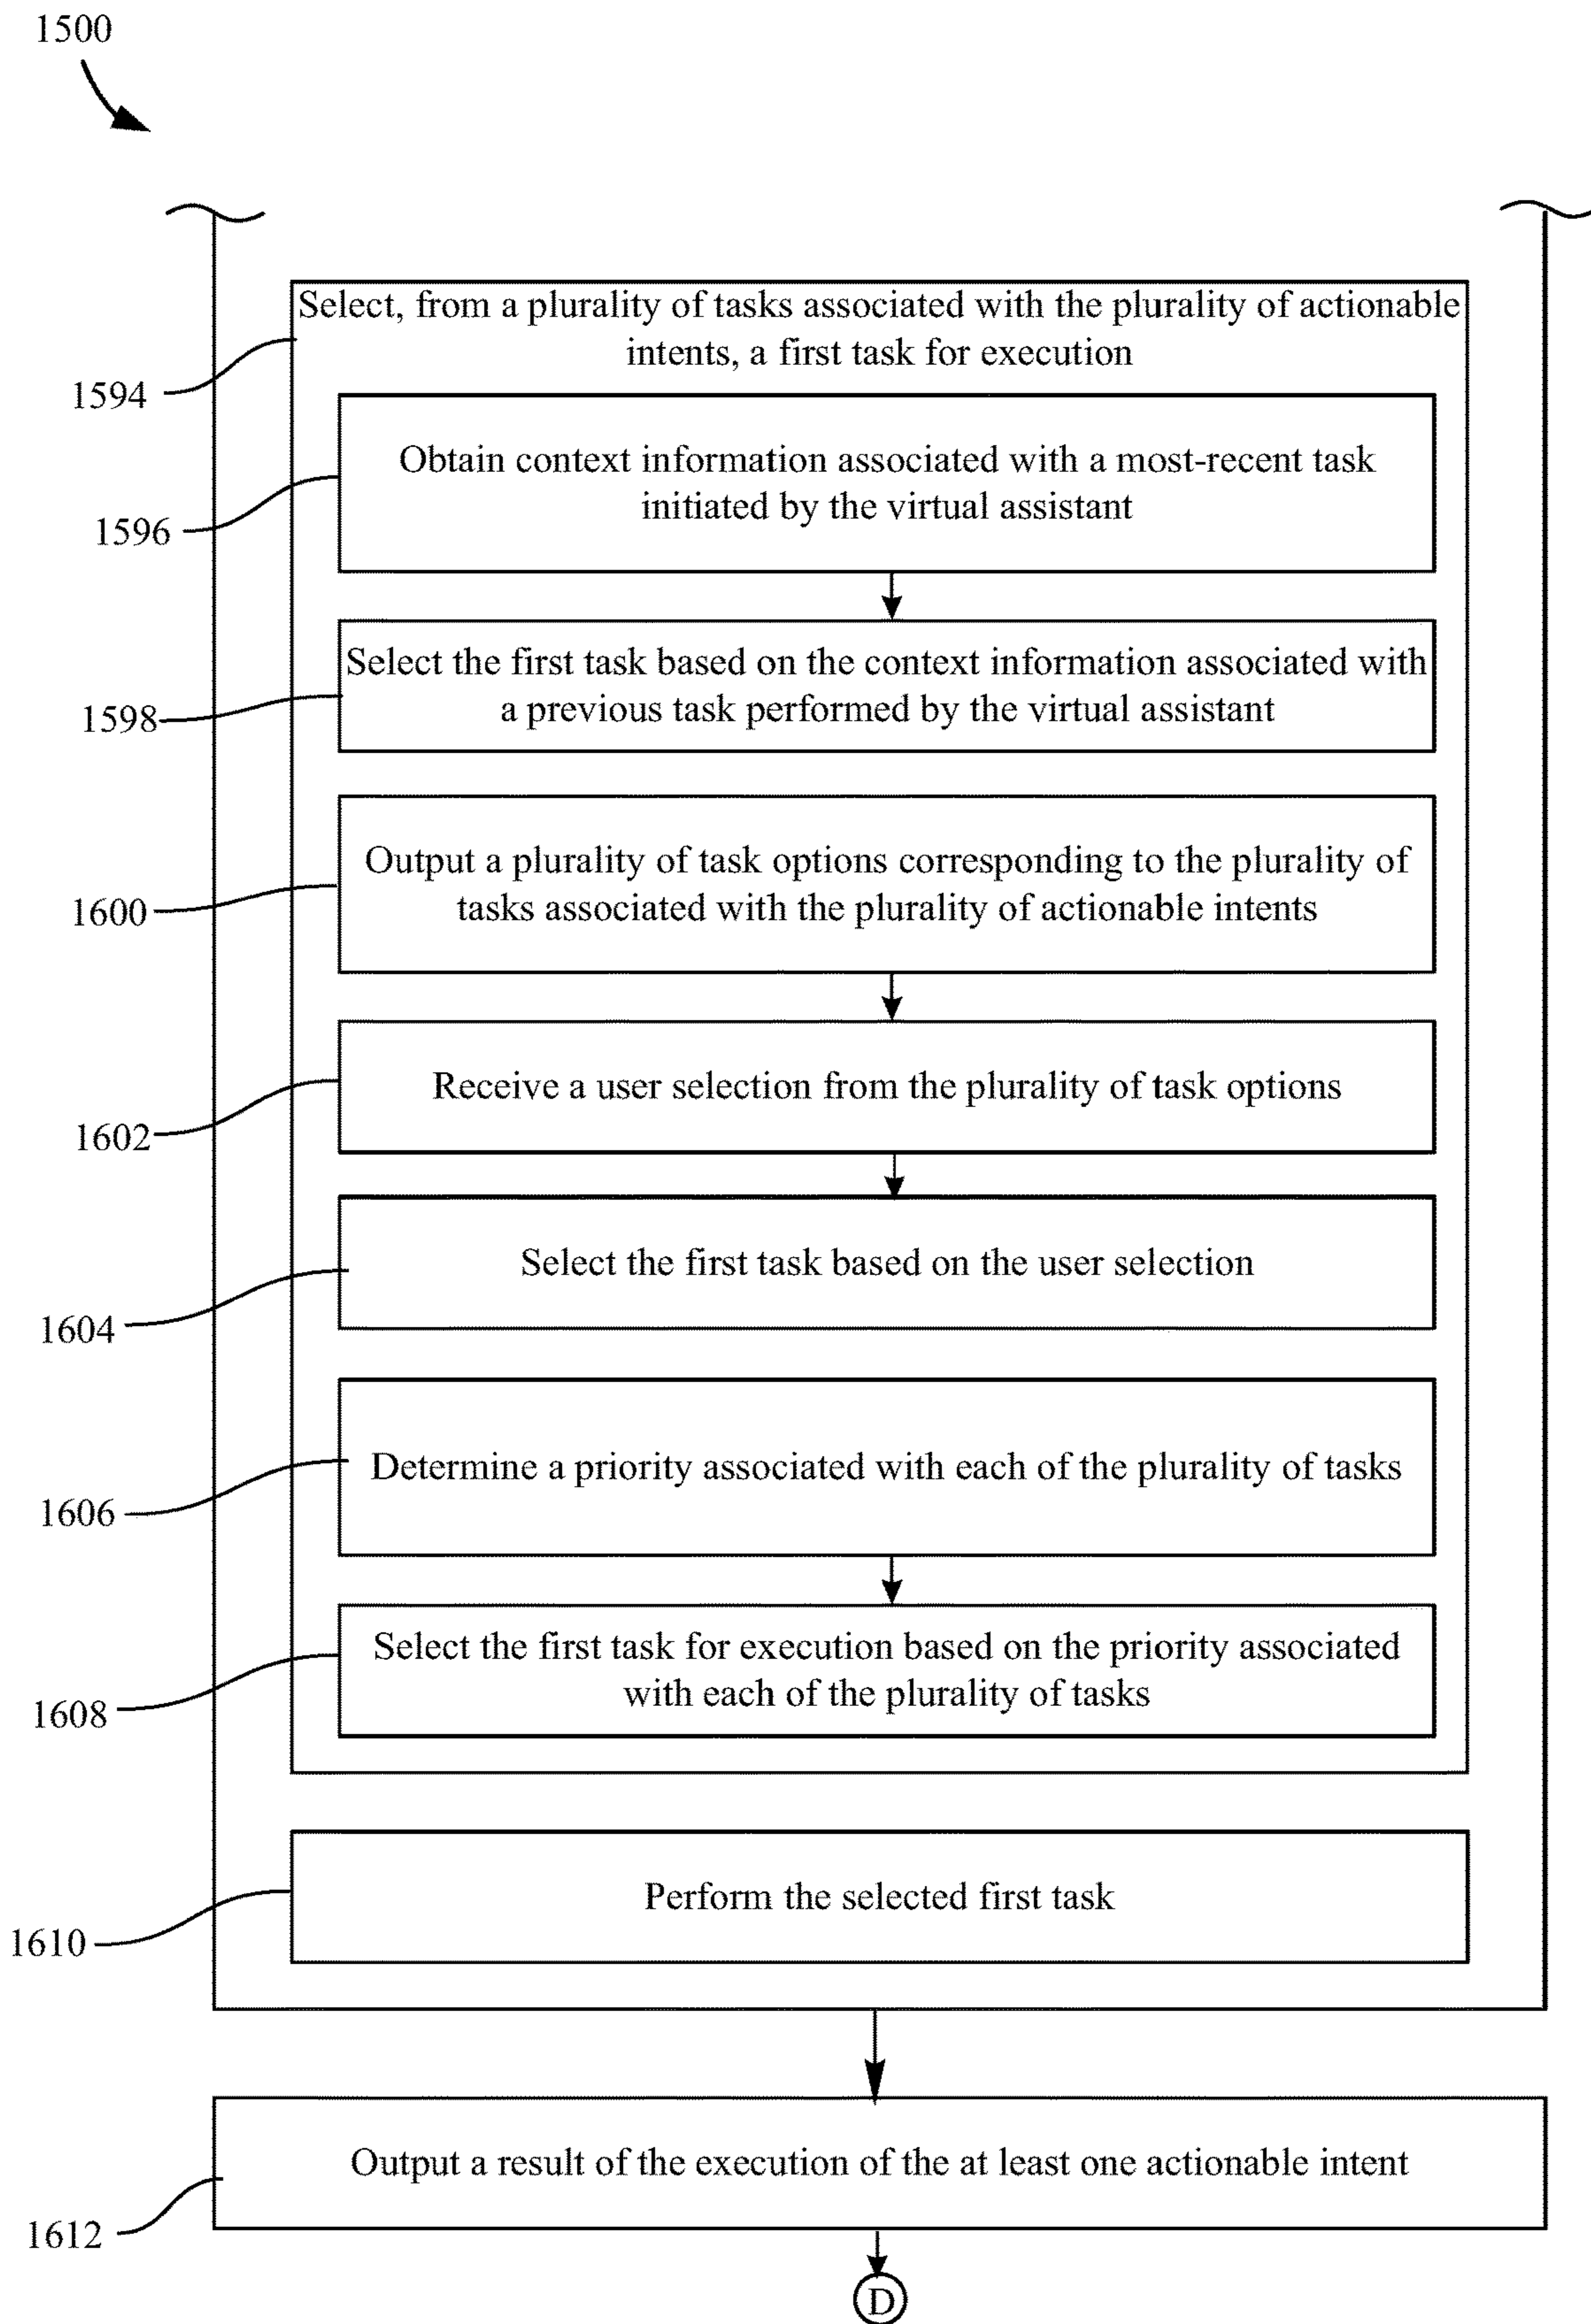


FIG. 15F

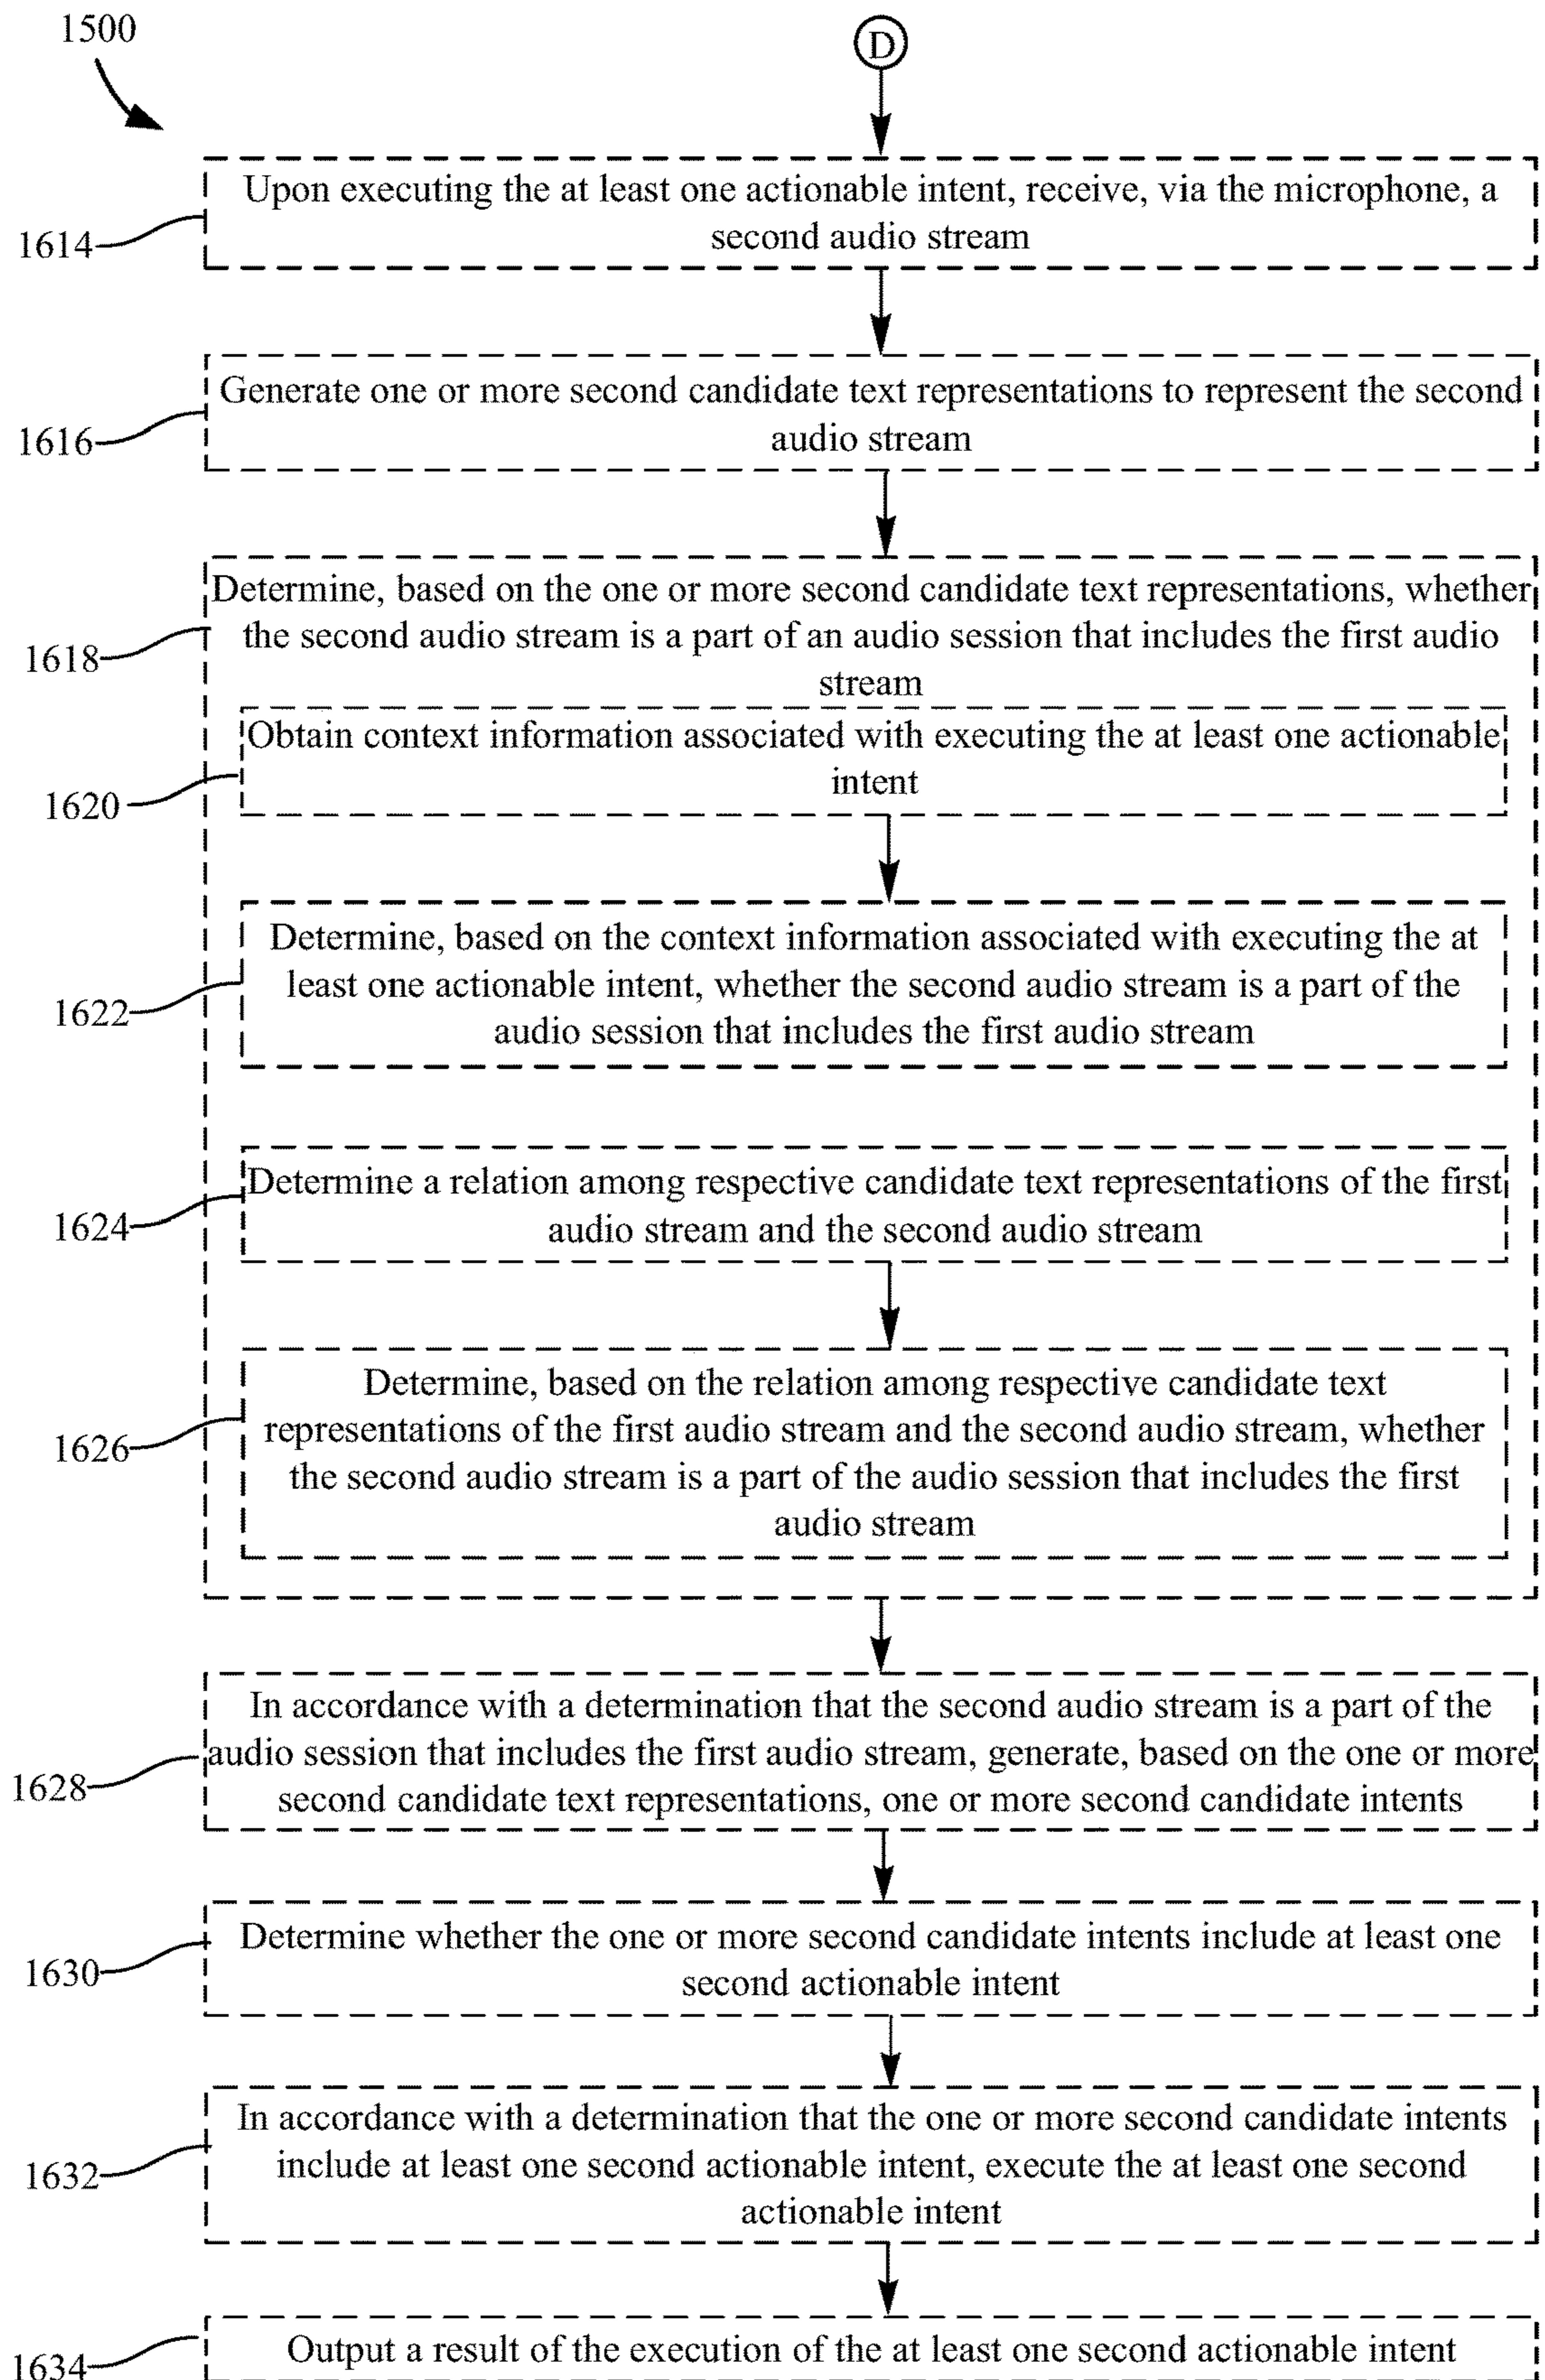


FIG. 15G

## NATURAL ASSISTANT INTERACTION

### CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of U.S. patent application Ser. No. 17/066,228, filed Oct. 8, 2020 entitled “NATURAL ASSISTANT INTERACTION,” which claims priority to U.S. patent application Ser. No. 16/019,331, now U.S. Pat. No. 10,818,288 filed Jun. 26, 2018, entitled “NATURAL ASSISTANT INTERACTION,” which claims priority to U.S. Provisional Application Ser. No. 62/648,084, entitled “NATURAL ASSISTANT INTERACTION,” filed on Mar. 26, 2018, the content of which are hereby incorporated by reference in their entirety.

### FIELD

[0002] This relates generally to virtual assistants and, more specifically, to providing natural language interaction by virtual assistants.

### BACKGROUND

[0003] Virtual assistants (or digital assistants or intelligent automated assistants) can provide a beneficial human-machine interface. Such assistants can allow users to interact with devices or systems using natural language in spoken and/or text forms. For example, a user can provide a speech input containing a user request to a digital assistant operating on an electronic device. The virtual assistant can interpret the user’s intent from the speech input and operationalize the user’s intent into tasks. The tasks can then be performed by executing one or more services of the electronic device, and a relevant output responsive to the user request can be returned to the user.

[0004] Virtual assistants can be activated upon receiving a trigger phrase such as “Hey Siri.” Upon activation, virtual assistants can receive and process user’s speech input. For example, a user’s speech input may include a leading trigger phrase to activate the virtual assistant followed by a request for information (e.g., “Hey Siri, how is the weather today?”). Leading every speech input with a trigger phrase (e.g., “Hey Siri”), however, can be inconvenient and quickly become cumbersome. It also does not represent a natural way of communication. For example, when a first user talks to a second user, the first user typically would not lead every sentence with the name of the second user. Thus, requiring the user to lead each speech input with a trigger phrase does not represent a natural way of communication and is less efficient.

### SUMMARY

[0005] Systems and processes for providing natural language interaction by a virtual assistant are provided.

[0006] In accordance with one or more examples, a method includes, at an electronic device with one or more processors, memory, and a microphone: receiving, via the microphone, a first audio stream including one or more utterances and determining whether the first audio stream includes a lexical trigger. In accordance with a determination that the first audio stream includes the lexical trigger, the method further includes generating one or more candidate text representations of the one or more utterances and determining whether at least one candidate text representation of the one or more candidate text representations is to

be disregarded by the virtual assistant. In accordance with a determination that at least one candidate text representation is to be disregarded by the virtual assistant, the method further includes generating one or more candidate intents based on candidate text representations of the one or more candidate text representations other than the to be disregarded at least one candidate text representation. The method further includes determining whether the one or more candidate intents include at least one actionable intent. In accordance with a determination that the one or more candidate intents include at least one actionable intent, the method further includes executing the at least one actionable intent and outputting a result of the execution of the at least one actionable intent.

[0007] Example non-transitory computer-readable media are disclosed herein. An example non-transitory computer-readable storage medium stores one or more programs. The one or more programs comprise instructions, which when executed by one or more processors of an electronic device, cause the electronic device to receive, via a microphone, a first audio stream including one or more utterances; determine whether the first audio stream includes a lexical trigger; in accordance with a determination that the first audio stream includes the lexical trigger, generate one or more candidate text representations of the one or more utterances; determine whether at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant; in accordance with a determination that at least one candidate text representation is to be disregarded by the virtual assistant, generate one or more candidate intents based on candidate text representations of the one or more candidate text representations other than the to be disregarded at least one candidate text representation; determine whether the one or more candidate intents include at least one actionable intent; in accordance with a determination that the one or more candidate intents include at least one actionable intent, execute the at least one actionable intent; and output a result of the execution of the at least one actionable intent.

[0008] Example electronic devices are disclosed herein. An example electronic device comprises one or more processors; a memory; and one or more programs, where the one or more programs are stored in the memory and configured to be executed by the one or more processors, the one or more programs including instructions for, receiving, via the microphone, a first audio stream including one or more utterances; determining whether the first audio stream includes a lexical trigger; in accordance with a determination that the first audio stream includes the lexical trigger, generating one or more candidate text representations of the one or more utterances; determining whether at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant; in accordance with a determination that at least one candidate text representation is to be disregarded by the virtual assistant, generating one or more candidate intents based on candidate text representations of the one or more candidate text representations other than the to be disregarded at least one candidate text representation; determining whether the one or more candidate intents include at least one actionable intent; in accordance with a determination that the one or more candidate intents include at least

one actionable intent, executing the at least one actionable intent; outputting a result of the execution of the at least one actionable intent.

**[0009]** An example electronic device comprises means for receiving, via the microphone, a first audio stream including one or more utterances; means for determining whether the first audio stream includes a lexical trigger; in accordance with a determination that the first audio stream includes the lexical trigger, means for generating one or more candidate text representations of the one or more utterances; means for determining whether at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant; in accordance with a determination that at least one candidate text representation is to be disregarded by the virtual assistant, means for generating one or more candidate intents based on candidate text representations of the one or more candidate text representations other than the to be disregarded at least one candidate text representation; means for determining whether the one or more candidate intents include at least one actionable intent; in accordance with a determination that the one or more candidate intents include at least one actionable intent, means for executing the at least one actionable intent; and means for outputting a result of the execution of the at least one actionable intent.

**[0010]** Current techniques facilitating speech-based human-machine interaction typically require using a trigger phrase at the beginning portion of an utterance from the user. As described above, this requirement can cause the human-machine interaction to become cumbersome and make the human-machine user interface less natural and efficient. Various techniques for providing natural language interaction described in this application eliminates or reduces the need of this requirement to lead every user utterance with a trigger phrase. Instead, a trigger word or phrase can be placed in any portion of an audio stream that may include one or more user utterances. Moreover, the techniques described in this application do not require using a trigger phrase that includes a plurality of words (e.g., “Hey Siri”). A single word (e.g., “Siri”) can be used to indicate that the audio stream including the user utterances is directed to the virtual assistant. This enables a more natural way of communication.

**[0011]** Furthermore, various techniques for facilitating speech-based human-machine interaction described in this application enhance the operability of the device and makes the user-device interface more efficient (e.g., by not requiring leading every user utterance with a trigger phrase) which, additionally, reduces power usage and improves battery life of the device by enabling the user to use the device more quickly and efficiently.

#### BRIEF DESCRIPTION OF THE DRAWINGS

**[0012]** FIG. 1 is a block diagram illustrating a system and environment for implementing a digital assistant, according to various examples.

**[0013]** FIG. 2A is a block diagram illustrating a portable multifunction device implementing the client-side portion of a digital assistant, according to various examples.

**[0014]** FIG. 2B is a block diagram illustrating exemplary components for event handling, according to various examples.

**[0015]** FIG. 3 illustrates a portable multifunction device implementing the client-side portion of a digital assistant, according to various examples.

**[0016]** FIG. 4 is a block diagram of an exemplary multifunction device with a display and a touch-sensitive surface, according to various examples.

**[0017]** FIG. 5A illustrates an exemplary user interface for a menu of applications on a portable multifunction device, according to various examples.

**[0018]** FIG. 5B illustrates an exemplary user interface for a multifunction device with a touch-sensitive surface that is separate from the display, according to various examples.

**[0019]** FIG. 6A illustrates a personal electronic device, according to various examples.

**[0020]** FIG. 6B is a block diagram illustrating a personal electronic device, according to various examples.

**[0021]** FIG. 7A is a block diagram illustrating a digital assistant system or a server portion thereof, according to various examples.

**[0022]** FIG. 7B illustrates the functions of the digital assistant shown in FIG. 7A, according to various examples.

**[0023]** FIG. 7C illustrates a portion of an ontology, according to various examples.

**[0024]** FIG. 8 illustrates a block diagram of an exemplary virtual assistant for providing natural language interaction.

**[0025]** FIG. 9 illustrates an exemplary user interface for providing natural language interaction by a virtual assistant.

**[0026]** FIG. 10 illustrates a block diagram of an exemplary virtual assistant for providing natural language interaction using context information.

**[0027]** FIG. 11A illustrates an exemplary user interface for providing natural language interaction by a virtual assistant using context information associated with a usage pattern.

**[0028]** FIG. 11B illustrates an exemplary user interface for providing natural language interaction by a virtual assistant using context information associated with sensory data.

**[0029]** FIGS. 12A-12D illustrate exemplary user interfaces for providing natural language interaction by a virtual assistant using context information associated with executing a previously determined actionable intent.

**[0030]** FIGS. 13A-13B illustrate exemplary user interfaces for providing natural language interaction by a virtual assistant using context information associated with a relation of user utterances or audio streams.

**[0031]** FIGS. 14A-14D illustrate exemplary user interfaces for selecting a task from a plurality of tasks using context information.

**[0032]** FIGS. 15A-15G illustrate a process for providing natural language interaction by a virtual assistant, according to various embodiments.

#### DETAILED DESCRIPTION

**[0033]** In the following description of embodiments, reference is made to the accompanying drawings in which are shown by way of illustration specific examples that can be practiced. It is to be understood that other examples can be used and structural changes can be made without departing from the scope of the various examples.

**[0034]** Various techniques for facilitating a more natural human-machine interaction are described. The techniques include reducing or eliminating the need for leading a user utterance with a trigger phrase and using a false-trigger mitigator to improve accuracy associated with determining whether a user utterance is directed to a virtual assistant. The

techniques also include performing candidate intent evaluation without actual execution (e.g., making a dry run), thereby determining whether a candidate intent is actionable. This determination avoids wasting processing power, user confusion, and thus improves operational efficiency of the device.

**[0035]** Although the following description uses terms “first,” “second,” etc. to describe various elements, these elements should not be limited by the terms. These terms are only used to distinguish one element from another. For example, a first input could be termed a second input, and, similarly, a second input could be termed a first input, without departing from the scope of the various described examples. The first input and the second input are both inputs and, in some cases, are separate and different inputs.

**[0036]** The terminology used in the description of the various described examples herein is for the purpose of describing particular examples only and is not intended to be limiting. As used in the description of the various described examples 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 “includes,” “including,” “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.

**[0037]** The term “if” may be construed to mean “when” or “upon” or “in response to determining” or “in response to detecting,” depending on the context. Similarly, the phrase “if it is determined” or “if [a stated condition or event] is detected” may be construed to mean “upon determining” or “in response to determining” or “upon detecting [the stated condition or event]” or “in response to detecting [the stated condition or event],” depending on the context.

## 1. System and Environment

**[0038]** FIG. 1 illustrates a block diagram of system **100** according to various examples. In some examples, system **100** implements a digital assistant. 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 performs 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.

**[0039]** 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 includes

a provision of the requested informational answer, a performance of the requested task, or a combination of the two. For example, a user asks the digital assistant a question, such as “Where am I right now?” Based on the user’s current location, the digital assistant answers, “You are in Central Park near the west gate.” The user also requests the performance of a task, for example, “Please invite my friends to my girlfriend’s birthday party next week.” In response, the digital assistant can acknowledge the request by saying “Yes, right away,” and then send a suitable calendar invite on behalf of the user to each of the user’s friends listed in the user’s electronic address book. 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.

**[0040]** As shown in FIG. 1, in some examples, a digital assistant is implemented according to a client-server model. The digital assistant includes client-side portion **102** (hereafter “DA client **102**”) executed on user device **104** and server-side portion **106** (hereafter “DA server **106**”) executed on server system **108**. DA client **102** communicates with DA server **106** through one or more networks **110**. DA client **102** provides client-side functionalities such as user-facing input and output processing and communication with DA server **106**. DA server **106** provides server-side functionalities for any number of DA clients **102** each residing on a respective user device **104**.

**[0041]** In some examples, DA server **106** includes client-facing I/O interface **112**, one or more processing modules **114**, data and models **116**, and I/O interface to external services **118**. The client-facing I/O interface **112** facilitates the client-facing input and output processing for DA server **106**. One or more processing modules **114** utilize data and models **116** to process speech input and determine the user’s intent based on natural language input. Further, one or more processing modules **114** perform task execution based on inferred user intent. In some examples, DA server **106** communicates with external services **120** through network (s) **110** for task completion or information acquisition. I/O interface to external services **118** facilitates such communications.

**[0042]** User device **104** can be any suitable electronic device. In some examples, user device is a portable multifunctional device (e.g., device **200**, described below with reference to FIG. 2A), a multifunctional device (e.g., device **400**, described below with reference to FIG. 4), or a personal electronic device (e.g., device **600**, described below with reference to FIG. 6A-B.) A portable multifunctional device is, for example, a mobile telephone that also contains other functions, such as PDA and/or music player functions. Specific examples of portable multifunction devices include the iPhone®, iPod Touch®, and iPad® devices from Apple Inc. of Cupertino, California. Other examples of portable multifunction devices include, without limitation, laptop or tablet computers. Further, in some examples, user device **104** is a non-portable multifunctional device. In particular, user device **104** is a desktop computer, a game console, a television, or a television set-top box. In some examples,

user device **104** includes a touch-sensitive surface (e.g., touch screen displays and/or touchpads). Further, user device **104** optionally includes one or more other physical user-interface devices, such as a physical keyboard, a mouse, and/or a joystick. Various examples of electronic devices, such as multifunctional devices, are described below in greater detail.

[0043] Examples of communication network(s) **110** include local area networks (LAN) and wide area networks (WAN), e.g., the Internet. Communication network(s) **110** is implemented using any known network protocol, including various wired or wireless protocols, such as, for example, Ethernet, Universal Serial Bus (USB), FIREWIRE, 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), WiMAX, or any other suitable communication protocol.

[0044] Server system **108** is implemented on one or more standalone data processing apparatus or a distributed network of computers. In some examples, 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 server system **108**.

[0045] In some examples, user device **104** communicates with DA server **106** via second user device **122**. Second user device **122** is similar or identical to user device **104**. For example, second user device **122** is similar to devices **200**, **400**, or **600** described below with reference to FIGS. 2A, 4, and 6A-B. User device **104** is configured to communicatively couple to second user device **122** via a direct communication connection, such as Bluetooth, NFC, BTLE, or the like, or via a wired or wireless network, such as a local Wi-Fi network. In some examples, second user device **122** is configured to act as a proxy between user device **104** and DA server **106**. For example, DA client **102** of user device **104** is configured to transmit information (e.g., a user request received at user device **104**) to DA server **106** via second user device **122**. DA server **106** processes the information and return relevant data (e.g., data content responsive to the user request) to user device **104** via second user device **122**.

[0046] In some examples, user device **104** is configured to communicate abbreviated requests for data to second user device **122** to reduce the amount of information transmitted from user device **104**. Second user device **122** is configured to determine supplemental information to add to the abbreviated request to generate a complete request to transmit to DA server **106**. This system architecture can advantageously allow user device **104** having limited communication capabilities and/or limited battery power (e.g., a watch or a similar compact electronic device) to access services provided by DA server **106** by using second user device **122**, having greater communication capabilities and/or battery power (e.g., a mobile phone, laptop computer, tablet computer, or the like), as a proxy to DA server **106**. While only two user devices **104** and **122** are shown in FIG. 1, it should be appreciated that system **100**, in some examples, includes any number and type of user devices configured in this proxy configuration to communicate with DA server system **106**.

[0047] Although the digital assistant shown in FIG. 1 includes both a client-side portion (e.g., DA client **102**) and a server-side portion (e.g., DA server **106**), in some

examples, the functions of a digital assistant are 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 implementations. For instance, in some examples, 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.

## 2. Electronic Devices

[0048] Attention is now directed toward embodiments of electronic devices for implementing the client-side portion of a digital assistant. FIG. 2A is a block diagram illustrating portable multifunction device **200** with touch-sensitive display system **212** in accordance with some embodiments. Touch-sensitive display **212** is sometimes called a “touch screen” for convenience and is sometimes known as or called a “touch-sensitive display system.” Device **200** includes memory **202** (which optionally includes one or more computer-readable storage mediums), memory controller **222**, one or more processing units (CPUs) **220**, peripherals interface **218**, RF circuitry **208**, audio circuitry **210**, speaker **211**, microphone **213**, input/output (I/O) subsystem **206**, other input control devices **216**, and external port **224**. Device **200** optionally includes one or more optical sensors **264**. Device **200** optionally includes one or more contact intensity sensors **265** for detecting intensity of contacts on device **200** (e.g., a touch-sensitive surface such as touch-sensitive display system **212** of device **200**). Device **200** optionally includes one or more tactile output generators **267** for generating tactile outputs on device **200** (e.g., generating tactile outputs on a touch-sensitive surface such as touch-sensitive display system **212** of device **200** or touchpad **455** of device **400**). These components optionally communicate over one or more communication buses or signal lines **203**.

[0049] As used in the specification and claims, the term “intensity” of a contact on a touch-sensitive surface refers to the force or pressure (force per unit area) of a contact (e.g., a finger contact) on the touch-sensitive surface, or to a substitute (proxy) for the force or pressure of a contact on the touch-sensitive surface. The intensity of a contact has a range of values that includes at least four distinct values and more typically includes hundreds of distinct values (e.g., at least 256). Intensity of a contact is, optionally, determined (or measured) using various approaches and various sensors or combinations of sensors. For example, one or more force sensors underneath or adjacent to the touch-sensitive surface are, optionally, used to measure force at various points on the touch-sensitive surface. In some implementations, force measurements from multiple force sensors are combined (e.g., a weighted average) to determine an estimated force of a contact. Similarly, a pressure-sensitive tip of a stylus is, optionally, used to determine a pressure of the stylus on the touch-sensitive surface. Alternatively, the size of the contact area detected on the touch-sensitive surface and/or changes thereto, the capacitance of the touch-sensitive surface proximate to the contact and/or changes thereto, and/or the resistance of the touch-sensitive surface proximate to the contact and/or changes thereto are, optionally, used as a substitute for the force or pressure of the contact on the touch-sensitive surface. In some implementations, the substitute measurements for contact force or pressure are used

directly to determine whether an intensity threshold has been exceeded (e.g., the intensity threshold is described in units corresponding to the substitute measurements). In some implementations, the substitute measurements for contact force or pressure are converted to an estimated force or pressure, and the estimated force or pressure is used to determine whether an intensity threshold has been exceeded (e.g., the intensity threshold is a pressure threshold measured in units of pressure). Using the intensity of a contact as an attribute of a user input allows for user access to additional device functionality that may otherwise not be accessible by the user on a reduced-size device with limited real estate for displaying affordances (e.g., on a touch-sensitive display) and/or receiving user input (e.g., via a touch-sensitive display, a touch-sensitive surface, or a physical/mechanical control such as a knob or a button).

**[0050]** As used in the specification and claims, the term “tactile output” refers to physical displacement of a device relative to a previous position of the device, physical displacement of a component (e.g., a touch-sensitive surface) of a device relative to another component (e.g., housing) of the device, or displacement of the component relative to a center of mass of the device that will be detected by a user with the user’s sense of touch. For example, in situations where the device or the component of the device is in contact with a surface of a user that is sensitive to touch (e.g., a finger, palm, or other part of a user’s hand), the tactile output generated by the physical displacement will be interpreted by the user as a tactile sensation corresponding to a perceived change in physical characteristics of the device or the component of the device. For example, movement of a touch-sensitive surface (e.g., a touch-sensitive display or trackpad) is, optionally, interpreted by the user as a “down click” or “up click” of a physical actuator button. In some cases, a user will feel a tactile sensation such as an “down click” or “up click” even when there is no movement of a physical actuator button associated with the touch-sensitive surface that is physically pressed (e.g., displaced) by the user’s movements. As another example, movement of the touch-sensitive surface is, optionally, interpreted or sensed by the user as “roughness” of the touch-sensitive surface, even when there is no change in smoothness of the touch-sensitive surface. While such interpretations of touch by a user will be subject to the individualized sensory perceptions of the user, there are many sensory perceptions of touch that are common to a large majority of users. Thus, when a tactile output is described as corresponding to a particular sensory perception of a user (e.g., an “up click,” a “down click,” “roughness”), unless otherwise stated, the generated tactile output corresponds to physical displacement of the device or a component thereof that will generate the described sensory perception for a typical (or average) user.

**[0051]** It should be appreciated that device **200** is only one example of a portable multifunction device, and that device **200** optionally has more or fewer components than shown, optionally combines two or more components, or optionally has a different configuration or arrangement of the components. The various components shown in FIG. 2A are implemented in hardware, software, or a combination of both hardware and software, including one or more signal processing and/or application-specific integrated circuits.

**[0052]** Memory **202** includes one or more computer-readable storage mediums. The computer-readable storage mediums are, for example, tangible and non-transitory. Memory

**202** includes high-speed random access memory and also includes non-volatile memory, such as one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices. Memory controller **222** controls access to memory **202** by other components of device **200**.

**[0053]** In some examples, a non-transitory computer-readable storage medium of memory **202** is used to store instructions (e.g., for performing aspects of processes described below) for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. In other examples, the instructions (e.g., for performing aspects of the processes described below) are stored on a non-transitory computer-readable storage medium (not shown) of the server system **108** or are divided between the non-transitory computer-readable storage medium of memory **202** and the non-transitory computer-readable storage medium of server system **108**.

**[0054]** Peripherals interface **218** is used to couple input and output peripherals of the device to CPU **220** and memory **202**. The one or more processors **220** run or execute various software programs and/or sets of instructions stored in memory **202** to perform various functions for device **200** and to process data. In some embodiments, peripherals interface **218**, CPU **220**, and memory controller **222** are implemented on a single chip, such as chip **204**. In some other embodiments, they are implemented on separate chips.

**[0055]** RF (radio frequency) circuitry **208** receives and sends RF signals, also called electromagnetic signals. RF circuitry **208** converts electrical signals to/from electromagnetic signals and communicates with communications networks and other communications devices via the electromagnetic signals. RF circuitry **208** optionally includes well-known circuitry for performing these functions, including but not limited to an antenna system, an RF transceiver, one or more amplifiers, a tuner, one or more oscillators, a digital signal processor, a CODEC chipset, a subscriber identity module (SIM) card, memory, and so forth. RF circuitry **208** optionally communicates with networks, such as the Internet, also referred to as the World Wide Web (WWW), 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 by wireless communication. The RF circuitry **208** optionally includes well-known circuitry for detecting near field communication (NFC) fields, such as by a short-range communication radio. The wireless communication optionally uses any of a plurality of communications standards, protocols, and technologies, including but not limited to Global System for Mobile Communications (GSM), Enhanced Data GSM Environment (EDGE), high-speed downlink packet access (HSDPA), high-speed uplink packet access (HSUPA), Evolution, Data-Only (EV-DO), HSPA, HSPA+, Dual-Cell HSPA (DC-HSPDA), long term evolution (LTE), near field communication (NFC), wideband code division multiple access (W-CDMA), code division multiple access (CDMA), time division multiple access (TDMA), Bluetooth, Bluetooth Low Energy (BTLE), Wireless Fidelity (Wi-Fi) (e.g., IEEE 802.11a, IEEE 802.11b, IEEE 802.11g, IEEE 802.11n, and/or IEEE 802.11ac), voice over Internet Protocol (VoIP), Wi-MAX, a protocol for e mail (e.g., Internet message



access protocol (IMAP) and/or post office protocol (POP)), instant messaging (e.g., extensible messaging and presence protocol (XMPP), Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE), Instant Messaging and Presence Service (IMPS)), and/or Short Message Service (SMS), or any other suitable communication protocol, including communication protocols not yet developed as of the filing date of this document.

**[0056]** Audio circuitry **210**, speaker **211**, and microphone **213** provide an audio interface between a user and device **200**. Audio circuitry **210** receives audio data from peripherals interface **218**, converts the audio data to an electrical signal, and transmits the electrical signal to speaker **211**. Speaker **211** converts the electrical signal to human-audible sound waves. Audio circuitry **210** also receives electrical signals converted by microphone **213** from sound waves. Audio circuitry **210** converts the electrical signal to audio data and transmits the audio data to peripherals interface **218** for processing. Audio data are retrieved from and/or transmitted to memory **202** and/or RF circuitry **208** by peripherals interface **218**. In some embodiments, audio circuitry **210** also includes a headset jack (e.g., **312**, FIG. 3). The headset jack provides an interface between audio circuitry **210** and removable audio input/output peripherals, such as output-only headphones or a headset with both output (e.g., a headphone for one or both ears) and input (e.g., a microphone).

**[0057]** I/O subsystem **206** couples input/output peripherals on device **200**, such as touch screen **212** and other input control devices **216**, to peripherals interface **218**. I/O subsystem **206** optionally includes display controller **256**, optical sensor controller **258**, intensity sensor controller **259**, haptic feedback controller **261**, and one or more input controllers **260** for other input or control devices. The one or more input controllers **260** receive/send electrical signals from/to other input control devices **216**. The other input control devices **216** optionally include physical buttons (e.g., push buttons, rocker buttons, etc.), dials, slider switches, joysticks, click wheels, and so forth. In some alternate embodiments, input controller(s) **260** are, optionally, coupled to any (or none) of the following: a keyboard, an infrared port, a USB port, and a pointer device such as a mouse. The one or more buttons (e.g., **308**, FIG. 3) optionally include an up/down button for volume control of speaker **211** and/or microphone **213**. The one or more buttons optionally include a push button (e.g., **306**, FIG. 3).

**[0058]** A quick press of the push button disengages a lock of touch screen **212** or begin a process that uses gestures on the touch screen to unlock the device, as described in U.S. patent application Ser. No. 11/322,549, “Unlocking a Device by Performing Gestures on an Unlock Image,” filed Dec. 23, 2005, U.S. Pat. No. 7,657,849, which is hereby incorporated by reference in its entirety. A longer press of the push button (e.g., **306**) turns power to device **200** on or off. The user is able to customize a functionality of one or more of the buttons. Touch screen **212** is used to implement virtual or soft buttons and one or more soft keyboards.

**[0059]** Touch-sensitive display **212** provides an input interface and an output interface between the device and a user. Display controller **256** receives and/or sends electrical signals from/to touch screen **212**. Touch screen **212** displays visual output to the user. The visual output includes graphics, text, icons, video, and any combination thereof (collec-

tively termed “graphics”). In some embodiments, some or all of the visual output correspond to user-interface objects.

**[0060]** Touch screen **212** has a touch-sensitive surface, sensor, or set of sensors that accepts input from the user based on haptic and/or tactile contact. Touch screen **212** and display controller **256** (along with any associated modules and/or sets of instructions in memory **202**) detect contact (and any movement or breaking of the contact) on touch screen **212** and convert the detected contact into interaction with user-interface objects (e.g., one or more soft keys, icons, web pages, or images) that are displayed on touch screen **212**. In an exemplary embodiment, a point of contact between touch screen **212** and the user corresponds to a finger of the user.

**[0061]** Touch screen **212** uses LCD (liquid crystal display) technology, LPD (light emitting polymer display) technology, or LED (light emitting diode) technology, although other display technologies may be used in other embodiments. Touch screen **212** and display controller **256** detect contact and any movement or breaking thereof using any of a plurality of touch sensing technologies now known or later developed, including but not limited to capacitive, resistive, infrared, and surface acoustic wave technologies, as well as other proximity sensor arrays or other elements for determining one or more points of contact with touch screen **212**. In an exemplary embodiment, projected mutual capacitance sensing technology is used, such as that found in the iPhone® and iPod Touch® from Apple Inc. of Cupertino, California.

**[0062]** A touch-sensitive display in some embodiments of touch screen **212** is analogous to the multi-touch sensitive touchpads described in the following U.S. Pat. No. 6,323,846 (Westerman et al.), U.S. Pat. No. 6,570,557 (Westerman et al.), and/or U.S. Pat. No. 6,677,932 (Westerman), and/or U.S. Patent Publication 2002/0015024A1, each of which is hereby incorporated by reference in its entirety. However, touch screen **212** displays visual output from device **200**, whereas touch-sensitive touchpads do not provide visual output.

**[0063]** A touch-sensitive display in some embodiments of touch screen **212** is as described in the following applications: (1) U.S. patent application Ser. No. 11/381,313, “Multipoint Touch Surface Controller,” filed May 2, 2006; (2) U.S. patent application Ser. No. 10/840,862, “Multipoint Touchscreen,” filed May 6, 2004; (3) U.S. patent application Ser. No. 10/903,964, “Gestures For Touch Sensitive Input Devices,” filed Jul. 30, 2004; (4) U.S. patent application Ser. No. 11/048,264, “Gestures For Touch Sensitive Input Devices,” filed Jan. 31, 2005; (5) U.S. patent application Ser. No. 11/038,590, “Mode-Based Graphical User Interfaces For Touch Sensitive Input Devices,” filed Jan. 18, 2005; (6) U.S. patent application Ser. No. 11/228,758, “Virtual Input Device Placement On A Touch Screen User Interface,” filed Sep. 16, 2005; (7) U.S. patent application Ser. No. 11/228,700, “Operation Of A Computer With A Touch Screen Interface,” filed Sep. 16, 2005; (8) U.S. patent application Ser. No. 11/228,737, “Activating Virtual Keys Of A Touch-Screen Virtual Keyboard,” filed Sep. 16, 2005; and (9) U.S. patent application Ser. No. 11/367,749, “Multi-Functional Hand-Held Device,” filed Mar. 3, 2006. All of these applications are incorporated by reference herein in their entirety.

**[0064]** Touch screen **212** has, for example, a video resolution in excess of 100 dpi. In some embodiments, the touch screen has a video resolution of approximately 160 dpi. The

user makes contact with touch screen **212** using any suitable object or appendage, such as a stylus, a finger, and so forth. In some embodiments, the user interface is designed to work primarily with finger-based contacts and gestures, which can be less precise than stylus-based input due to the larger area of contact of a finger on the touch screen. In some embodiments, the device translates the rough finger-based input into a precise pointer/cursor position or command for performing the actions desired by the user.

[0065] In some embodiments, in addition to the touch screen, device **200** includes a touchpad (not shown) for activating or deactivating particular functions. In some embodiments, the touchpad is a touch-sensitive area of the device that, unlike the touch screen, does not display visual output. The touchpad is a touch-sensitive surface that is separate from touch screen **212** or an extension of the touch-sensitive surface formed by the touch screen.

[0066] Device **200** also includes power system **262** for powering the various components. Power system **262** includes a power management system, one or more power sources (e.g., battery, alternating current (AC)), a recharging system, a power failure detection circuit, a power converter or inverter, a power status indicator (e.g., a light-emitting diode (LED)) and any other components associated with the generation, management and distribution of power in portable devices.

[0067] Device **200** also includes one or more optical sensors **264**. FIG. 2A shows an optical sensor coupled to optical sensor controller **258** in I/O subsystem **206**. Optical sensor **264** includes charge-coupled device (CCD) or complementary metal-oxide semiconductor (CMOS) phototransistors. Optical sensor **264** receives light from the environment, projected through one or more lenses, and converts the light to data representing an image. In conjunction with imaging module **243** (also called a camera module), optical sensor **264** captures still images or video. In some embodiments, an optical sensor is located on the back of device **200**, opposite touch screen display **212** on the front of the device so that the touch screen display is used as a viewfinder for still and/or video image acquisition. In some embodiments, an optical sensor is located on the front of the device so that the user's image is obtained for video conferencing while the user views the other video conference participants on the touch screen display. In some embodiments, the position of optical sensor **264** can be changed by the user (e.g., by rotating the lens and the sensor in the device housing) so that a single optical sensor **264** is used along with the touch screen display for both video conferencing and still and/or video image acquisition.

[0068] Device **200** optionally also includes one or more contact intensity sensors **265**. FIG. 2A shows a contact intensity sensor coupled to intensity sensor controller **259** in I/O subsystem **206**. Contact intensity sensor **265** optionally includes one or more piezoresistive strain gauges, capacitive force sensors, electric force sensors, piezoelectric force sensors, optical force sensors, capacitive touch-sensitive surfaces, or other intensity sensors (e.g., sensors used to measure the force (or pressure) of a contact on a touch-sensitive surface). Contact intensity sensor **265** receives contact intensity information (e.g., pressure information or a proxy for pressure information) from the environment. In some embodiments, at least one contact intensity sensor is collocated with, or proximate to, a touch-sensitive surface (e.g., touch-sensitive display system **212**). In some embodi-

ments, at least one contact intensity sensor is located on the back of device **200**, opposite touch screen display **212**, which is located on the front of device **200**.

[0069] Device **200** also includes one or more proximity sensors **266**. FIG. 2A shows proximity sensor **266** coupled to peripherals interface **218**. Alternately, proximity sensor **266** is coupled to input controller **260** in I/O subsystem **206**. Proximity sensor **266** is performed as described in U.S. patent application Ser. No. 11/241,839, "Proximity Detector In Handheld Device"; Ser. No. 11/240,788, "Proximity Detector In Handheld Device"; Ser. No. 11/620,702, "Using Ambient Light Sensor To Augment Proximity Sensor Output"; Ser. No. 11/586,862, "Automated Response To And Sensing Of User Activity In Portable Devices"; and Ser. No. 11/638,251, "Methods And Systems For Automatic Configuration Of Peripherals," which are hereby incorporated by reference in their entirety. In some embodiments, the proximity sensor turns off and disables touch screen **212** when the multifunction device is placed near the user's ear (e.g., when the user is making a phone call).

[0070] Device **200** optionally also includes one or more tactile output generators **267**. FIG. 2A shows a tactile output generator coupled to haptic feedback controller **261** in I/O subsystem **206**. Tactile output generator **267** optionally includes one or more electroacoustic devices such as speakers or other audio components and/or electromechanical devices that convert energy into linear motion such as a motor, solenoid, electroactive polymer, piezoelectric actuator, electrostatic actuator, or other tactile output generating component (e.g., a component that converts electrical signals into tactile outputs on the device). Contact intensity sensor **265** receives tactile feedback generation instructions from haptic feedback module **233** and generates tactile outputs on device **200** that are capable of being sensed by a user of device **200**. In some embodiments, at least one tactile output generator is collocated with, or proximate to, a touch-sensitive surface (e.g., touch-sensitive display system **212**) and, optionally, generates a tactile output by moving the touch-sensitive surface vertically (e.g., in/out of a surface of device **200**) or laterally (e.g., back and forth in the same plane as a surface of device **200**). In some embodiments, at least one tactile output generator sensor is located on the back of device **200**, opposite touch screen display **212**, which is located on the front of device **200**.

[0071] Device **200** also includes one or more accelerometers **268**. FIG. 2A shows accelerometer **268** coupled to peripherals interface **218**. Alternately, accelerometer **268** is coupled to an input controller **260** in I/O subsystem **206**. Accelerometer **268** performs, for example, as described in U.S. Patent Publication No. 20050190059, "Acceleration-based Theft Detection System for Portable Electronic Devices," and U.S. Patent Publication No. 20060017692, "Methods And Apparatuses For Operating A Portable Device Based On An Accelerometer," both of which are incorporated by reference herein in their entirety. In some embodiments, information is displayed on the touch screen display in a portrait view or a landscape view based on an analysis of data received from the one or more accelerometers. Device **200** optionally includes, in addition to accelerometer(s) **268**, a magnetometer (not shown) and a GPS (or GLONASS or other global navigation system) receiver (not shown) for obtaining information concerning the location and orientation (e.g., portrait or landscape) of device **200**.

[0072] In some embodiments, the software components stored in memory 202 include operating system 226, communication module (or set of instructions) 228, contact/motion module (or set of instructions) 230, graphics module (or set of instructions) 232, text input module (or set of instructions) 234, Global Positioning System (GPS) module (or set of instructions) 235, Digital Assistant Client Module 229, and applications (or sets of instructions) 236. Further, memory 202 stores data and models, such as user data and models 231. Furthermore, in some embodiments, memory 202 (FIG. 2A) or 470 (FIG. 4) stores device/global internal state 257, as shown in FIGS. 2A and 4. Device/global internal state 257 includes one or more of: active application state, indicating which applications, if any, are currently active; display state, indicating what applications, views or other information occupy various regions of touch screen display 212; sensor state, including information obtained from the device's various sensors and input control devices 216; and location information concerning the device's location and/or attitude.

[0073] Operating system 226 (e.g., Darwin, RTXC, LINUX, UNIX, OS X, iOS, 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 communication between various hardware and software components.

[0074] Communication module 228 facilitates communication with other devices over one or more external ports 224 and also includes various software components for handling data received by RF circuitry 208 and/or external port 224. External port 224 (e.g., Universal Serial Bus (USB), FIREWIRE, etc.) is adapted for coupling directly to other devices or indirectly over a network (e.g., the Internet, wireless LAN, etc.). In some embodiments, the external port is a multi-pin (e.g., 30-pin) connector that is the same as, or similar to and/or compatible with, the 30-pin connector used on iPod® (trademark of Apple Inc.) devices.

[0075] Contact/motion module 230 optionally detects contact with touch screen 212 (in conjunction with display controller 256) and other touch-sensitive devices (e.g., a touchpad or physical click wheel). Contact/motion module 230 includes various software components for performing various operations related to detection of contact, such as determining if contact has occurred (e.g., detecting a finger-down event), determining an intensity of the contact (e.g., the force or pressure of the contact or a substitute for the force or pressure of the contact), determining if there is movement of the contact and tracking the movement across the touch-sensitive surface (e.g., detecting one or more finger-dragging events), and determining if the contact has ceased (e.g., detecting a finger-up event or a break in contact). Contact/motion module 230 receives contact data from the touch-sensitive surface. Determining movement of the point of contact, which is represented by a series of contact data, optionally includes determining speed (magnitude), velocity (magnitude and direction), and/or an acceleration (a change in magnitude and/or direction) of the point of contact. These operations are, optionally, applied to single contacts (e.g., one finger contacts) or to multiple simultaneous contacts (e.g., "multitouch"/multiple finger contacts). In some embodiments, contact/motion module 230 and display controller 256 detect contact on a touchpad.

[0076] In some embodiments, contact/motion module 230 uses a set of one or more intensity thresholds to determine whether an operation has been performed by a user (e.g., to determine whether a user has "clicked" on an icon). In some embodiments, at least a subset of the intensity thresholds are determined in accordance with software parameters (e.g., the intensity thresholds are not determined by the activation thresholds of particular physical actuators and can be adjusted without changing the physical hardware of device 200). For example, a mouse "click" threshold of a trackpad or touch screen display can be set to any of a large range of predefined threshold values without changing the trackpad or touch screen display hardware. Additionally, in some implementations, a user of the device is provided with software settings for adjusting one or more of the set of intensity thresholds (e.g., by adjusting individual intensity thresholds and/or by adjusting a plurality of intensity thresholds at once with a system-level click "intensity" parameter).

[0077] Contact/motion module 230 optionally detects a gesture input by a user. Different gestures on the touch-sensitive surface have different contact patterns (e.g., different motions, timings, and/or intensities of detected contacts). Thus, a gesture is, optionally, detected by detecting a particular contact pattern. For example, detecting a finger tap gesture includes detecting a finger-down event followed by detecting a finger-up (liftoff) event at the same position (or substantially the same position) as the finger-down event (e.g., at the position of an icon). As another example, detecting a finger swipe gesture on the touch-sensitive surface includes detecting a finger-down event followed by detecting one or more finger-dragging events, and subsequently followed by detecting a finger-up (liftoff) event.

[0078] Graphics module 232 includes various known software components for rendering and displaying graphics on touch screen 212 or other display, including components for changing the visual impact (e.g., brightness, transparency, saturation, contrast, or other visual property) of graphics that are displayed. As used herein, the term "graphics" includes any object that can be displayed to a user, including, without limitation, text, web pages, icons (such as user-interface objects including soft keys), digital images, videos, animations, and the like.

[0079] In some embodiments, graphics module 232 stores data representing graphics to be used. Each graphic is, optionally, assigned a corresponding code. Graphics module 232 receives, from applications etc., one or more codes specifying graphics to be displayed along with, if necessary, coordinate data and other graphic property data, and then generates screen image data to output to display controller 256.

[0080] Haptic feedback module 233 includes various software components for generating instructions used by tactile output generator(s) 267 to produce tactile outputs at one or more locations on device 200 in response to user interactions with device 200.

[0081] Text input module 234, which is, in some examples, a component of graphics module 232, provides soft keyboards for entering text in various applications (e.g., contacts 237, email 240, IM 241, browser 247, and any other application that needs text input).

[0082] GPS module 235 determines the location of the device and provides this information for use in various applications (e.g., to telephone 238 for use in location-based dialing; to camera 243 as picture/video metadata; and to

applications that provide location-based services such as weather widgets, local yellow page widgets, and map/navigation widgets).

[0083] Digital assistant client module 229 includes various client-side digital assistant instructions to provide the client-side functionalities of the digital assistant. For example, digital assistant client module 229 is capable of accepting voice input (e.g., speech input), text input, touch input, and/or gestural input through various user interfaces (e.g., microphone 213, accelerometer(s) 268, touch-sensitive display system 212, optical sensor(s) 229, other input control devices 216, etc.) of portable multifunction device 200. Digital assistant client module 229 is also capable of providing output in audio (e.g., speech output), visual, and/or tactile forms through various output interfaces (e.g., speaker 211, touch-sensitive display system 212, tactile output generator(s) 267, etc.) of portable multifunction device 200. For example, output is provided as voice, sound, alerts, text messages, menus, graphics, videos, animations, vibrations, and/or combinations of two or more of the above. During operation, digital assistant client module 229 communicates with DA server 106 using RF circuitry 208.

[0084] User data and models 231 include various data associated with the user (e.g., user-specific vocabulary data, user preference data, user-specified name pronunciations, data from the user's electronic address book, to-do lists, shopping lists, etc.) to provide the client-side functionalities of the digital assistant. Further, user data and models 231 include various models (e.g., speech recognition models, statistical language models, natural language processing models, ontology, task flow models, service models, etc.) for processing user input and determining user intent.

[0085] In some examples, digital assistant client module 229 utilizes the various sensors, subsystems, and peripheral devices of portable multifunction device 200 to gather additional information from the surrounding environment of the portable multifunction device 200 to establish a context associated with a user, the current user interaction, and/or the current user input. In some examples, digital assistant client module 229 provides the contextual information or a subset thereof with the user input to DA server 106 to help infer the user's intent. In some examples, the digital assistant also uses the contextual information to determine how to prepare and deliver outputs to the user. Contextual information is referred to as context data.

[0086] In some examples, the contextual 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 examples, the contextual information can also include the physical state of the device, e.g., device orientation, device location, device temperature, power level, speed, acceleration, motion patterns, cellular signals strength, etc. In some examples, information related to the software state of DA server 106, e.g., running processes, installed programs, past and present network activities, background services, error logs, resources usage, etc., and of portable multifunction device 200 is provided to DA server 106 as contextual information associated with a user input.

[0087] In some examples, the digital assistant client module 229 selectively provides information (e.g., user data 231) stored on the portable multifunction device 200 in response to requests from DA server 106. In some examples, digital assistant client module 229 also elicits additional input from

the user via a natural language dialogue or other user interfaces upon request by DA server 106. Digital assistant client module 229 passes the additional input to DA server 106 to help DA server 106 in intent deduction and/or fulfillment of the user's intent expressed in the user request.

[0088] A more detailed description of a digital assistant is described below with reference to FIGS. 7A-C. It should be recognized that digital assistant client module 229 can include any number of the sub-modules of digital assistant module 726 described below.

[0089] Applications 236 include the following modules (or sets of instructions), or a subset or superset thereof:

- [0090] Contacts module 237 (sometimes called an address book or contact list);
- [0091] Telephone module 238;
- [0092] Video conference module 239;
- [0093] E-mail client module 240;
- [0094] Instant messaging (IM) module 241;
- [0095] Workout support module 242;
- [0096] Camera module 243 for still and/or video images;
- [0097] Image management module 244;
- [0098] Video player module;
- [0099] Music player module;
- [0100] Browser module 247;
- [0101] Calendar module 248;
- [0102] Widget modules 249, which includes, in some examples, one or more of: weather widget 249-1, stocks widget 249-2, calculator widget 249-3, alarm clock widget 249-4, dictionary widget 249-5, and other widgets obtained by the user, as well as user-created widgets 249-6;
- [0103] Widget creator module 250 for making user-created widgets 249-6;
- [0104] Search module 251;
- [0105] Video and music player module 252, which merges video player module and music player module;
- [0106] Notes module 253;
- [0107] Map module 254; and/or
- [0108] Online video module 255.

[0109] Examples of other applications 236 that are stored in memory 202 include other word processing applications, other image editing applications, drawing applications, presentation applications, JAVA-enabled applications, encryption, digital rights management, voice recognition, and voice replication.

[0110] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, contacts module 237 are used to manage an address book or contact list (e.g., stored in application internal state 292 of contacts module 237 in memory 202 or memory 470), including: adding name(s) to the address book; deleting name(s) from the address book; associating telephone number(s), e-mail address(es), physical address(es) or other information with a name; associating an image with a name; categorizing and sorting names; providing telephone numbers or e-mail addresses to initiate and/or facilitate communications by telephone 238, video conference module 239, e-mail 240, or IM 241; and so forth.

[0111] In conjunction with RF circuitry 208, audio circuitry 210, speaker 211, microphone 213, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, telephone module 238 are used to enter a sequence of characters corresponding

to a telephone number, access one or more telephone numbers in contacts module 237, modify a telephone number that has been entered, dial a respective telephone number, conduct a conversation, and disconnect or hang up when the conversation is completed. As noted above, the wireless communication uses any of a plurality of communications standards, protocols, and technologies.

[0112] In conjunction with RF circuitry 208, audio circuitry 210, speaker 211, microphone 213, touch screen 212, display controller 256, optical sensor 264, optical sensor controller 258, contact/motion module 230, graphics module 232, text input module 234, contacts module 237, and telephone module 238, video conference module 239 includes executable instructions to initiate, conduct, and terminate a video conference between a user and one or more other participants in accordance with user instructions.

[0113] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, e-mail client module 240 includes executable instructions to create, send, receive, and manage e-mail in response to user instructions. In conjunction with image management module 244, e-mail client module 240 makes it very easy to create and send e-mails with still or video images taken with camera module 243.

[0114] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, the instant messaging module 241 includes executable instructions to enter a sequence of characters corresponding to an instant message, to modify previously entered characters, to transmit a respective instant message (for example, using a Short Message Service (SMS) or Multimedia Message Service (MMS) protocol for telephony-based instant messages or using XMPP, SIMPLE, or IMPS for Internet-based instant messages), to receive instant messages, and to view received instant messages. In some embodiments, transmitted and/or received instant messages include graphics, photos, audio files, video files and/or other attachments as are supported in an MMS and/or an Enhanced Messaging Service (EMS). As used herein, “instant messaging” refers to both telephony-based messages (e.g., messages sent using SMS or MMS) and Internet-based messages (e.g., messages sent using XMPP, SIMPLE, or IMPS).

[0115] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, GPS module 235, map module 254, and music player module, workout support module 242 includes executable instructions to create workouts (e.g., with time, distance, and/or calorie burning goals); communicate with workout sensors (sports devices); receive workout sensor data; calibrate sensors used to monitor a workout; select and play music for a workout; and display, store, and transmit workout data.

[0116] In conjunction with touch screen 212, display controller 256, optical sensor(s) 264, optical sensor controller 258, contact/motion module 230, graphics module 232, and image management module 244, camera module 243 includes executable instructions to capture still images or video (including a video stream) and store them into memory 202, modify characteristics of a still image or video, or delete a still image or video from memory 202.

[0117] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module

232, text input module 234, and camera module 243, image management module 244 includes executable instructions to arrange, modify (e.g., edit), or otherwise manipulate, label, delete, present (e.g., in a digital slide show or album), and store still and/or video images.

[0118] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, browser module 247 includes executable instructions to browse the Internet in accordance with user instructions, including searching, linking to, receiving, and displaying web pages or portions thereof, as well as attachments and other files linked to web pages.

[0119] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, e-mail client module 240, and browser module 247, calendar module 248 includes executable instructions to create, display, modify, and store calendars and data associated with calendars (e.g., calendar entries, to-do lists, etc.) in accordance with user instructions.

[0120] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and browser module 247, widget modules 249 are mini-applications that can be downloaded and used by a user (e.g., weather widget 249-1, stocks widget 249-2, calculator widget 249-3, alarm clock widget 249-4, and dictionary widget 249-5) or created by the user (e.g., user-created widget 249-6). In some embodiments, a widget includes an HTML (Hypertext Markup Language) file, a CSS (Cascading Style Sheets) file, and a JavaScript file. In some embodiments, a widget includes an XML (Extensible Markup Language) file and a JavaScript file (e.g., Yahoo! Widgets).

[0121] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, and browser module 247, the widget creator module 250 are used by a user to create widgets (e.g., turning a user-specified portion of a web page into a widget).

[0122] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, search module 251 includes executable instructions to search for text, music, sound, image, video, and/or other files in memory 202 that match one or more search criteria (e.g., one or more user-specified search terms) in accordance with user instructions.

[0123] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, audio circuitry 210, speaker 211, RF circuitry 208, and browser module 247, video and music player module 252 includes executable instructions that allow the user to download and play back recorded music and other sound files stored in one or more file formats, such as MP3 or AAC files, and executable instructions to display, present, or otherwise play back videos (e.g., on touch screen 212 or on an external, connected display via external port 224). In some embodiments, device 200 optionally includes the functionality of an MP3 player, such as an iPod (trademark of Apple Inc.).

[0124] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, and text input module 234, notes module 253 includes executable instructions to create and manage notes, to-do lists, and the like in accordance with user instructions.

[0125] In conjunction with RF circuitry 208, touch screen 212, display controller 256, contact/motion module 230, graphics module 232, text input module 234, GPS module 235, and browser module 247, map module 254 are used to receive, display, modify, and store maps and data associated with maps (e.g., driving directions, data on stores and other points of interest at or near a particular location, and other location-based data) in accordance with user instructions.

[0126] In conjunction with touch screen 212, display controller 256, contact/motion module 230, graphics module 232, audio circuitry 210, speaker 211, RF circuitry 208, text input module 234, e-mail client module 240, and browser module 247, online video module 255 includes instructions that allow the user to access, browse, receive (e.g., by streaming and/or download), play back (e.g., on the touch screen or on an external, connected display via external port 224), send an e-mail with a link to a particular online video, and otherwise manage online videos in one or more file formats, such as H.264. In some embodiments, instant messaging module 241, rather than e-mail client module 240, is used to send a link to a particular online video. Additional description of the online video application can be found in U.S. Provisional Patent Application No. 60/936,562, “Portable Multifunction Device, Method, and Graphical User Interface for Playing Online Videos,” filed Jun. 20, 2007, and U.S. patent application Ser. No. 11/968,067, “Portable Multifunction Device, Method, and Graphical User Interface for Playing Online Videos,” filed Dec. 31, 2007, the contents of which are hereby incorporated by reference in their entirety.

[0127] Each of the above-identified modules and applications corresponds to a set of executable instructions for performing one or more functions described above and the methods described in this application (e.g., the computer-implemented methods and other information processing methods described herein). These modules (e.g., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules can be combined or otherwise rearranged in various embodiments. For example, video player module can be combined with music player module into a single module (e.g., video and music player module 252, FIG. 2A). In some embodiments, memory 202 stores a subset of the modules and data structures identified above. Furthermore, memory 202 stores additional modules and data structures not described above.

[0128] In some embodiments, device 200 is a device where operation of a predefined set of functions on the device is performed exclusively through a touch screen and/or a touchpad. By using a touch screen and/or a touchpad as the primary input control device for operation of device 200, the number of physical input control devices (such as push buttons, dials, and the like) on device 200 is reduced.

[0129] The predefined set of functions that are performed exclusively through a touch screen and/or a touchpad optionally include navigation between user interfaces. In some embodiments, the touchpad, when touched by the user, navigates device 200 to a main, home, or root menu from any user interface that is displayed on device 200. In such embodiments, a “menu button” is implemented using a touchpad. In some other embodiments, the menu button is a physical push button or other physical input control device instead of a touchpad.

[0130] FIG. 2B is a block diagram illustrating exemplary components for event handling in accordance with some embodiments. In some embodiments, memory 202 (FIG. 2A) or 470 (FIG. 4) includes event sorter 270 (e.g., in operating system 226) and a respective application 236-1 (e.g., any of the aforementioned applications 237-251, 255, 480-490).

[0131] Event sorter 270 receives event information and determines the application 236-1 and application view 291 of application 236-1 to which to deliver the event information. Event sorter 270 includes event monitor 271 and event dispatcher module 274. In some embodiments, application 236-1 includes application internal state 292, which indicates the current application view(s) displayed on touch-sensitive display 212 when the application is active or executing. In some embodiments, device/global internal state 257 is used by event sorter 270 to determine which application(s) is (are) currently active, and application internal state 292 is used by event sorter 270 to determine application views 291 to which to deliver event information.

[0132] In some embodiments, application internal state 292 includes additional information, such as one or more of: resume information to be used when application 236-1 resumes execution, user interface state information that indicates information being displayed or that is ready for display by application 236-1, a state queue for enabling the user to go back to a prior state or view of application 236-1, and a redo/undo queue of previous actions taken by the user.

[0133] Event monitor 271 receives event information from peripherals interface 218. Event information includes information about a sub-event (e.g., a user touch on touch-sensitive display 212, as part of a multi-touch gesture). Peripherals interface 218 transmits information it receives from I/O subsystem 206 or a sensor, such as proximity sensor 266, accelerometer(s) 268, and/or microphone 213 (through audio circuitry 210). Information that peripherals interface 218 receives from I/O subsystem 206 includes information from touch-sensitive display 212 or a touch-sensitive surface.

[0134] In some embodiments, event monitor 271 sends requests to the peripherals interface 218 at predetermined intervals. In response, peripherals interface 218 transmits event information. In other embodiments, peripherals interface 218 transmits event information only when there is a significant event (e.g., receiving an input above a predetermined noise threshold and/or for more than a predetermined duration).

[0135] In some embodiments, event sorter 270 also includes a hit view determination module 272 and/or an active event recognizer determination module 273.

[0136] Hit view determination module 272 provides software procedures for determining where a sub-event has taken place within one or more views when touch-sensitive display 212 displays more than one view. Views are made up of controls and other elements that a user can see on the display.

[0137] Another aspect of the user interface associated with an application is a set of views, sometimes herein called application views or user interface windows, in which information is displayed and touch-based gestures occur. The application views (of a respective application) in which a touch is detected correspond to programmatic levels within a programmatic or view hierarchy of the application. For example, the lowest level view in which a touch is detected

is called the hit view, and the set of events that are recognized as proper inputs is determined based, at least in part, on the hit view of the initial touch that begins a touch-based gesture.

[0138] Hit view determination module 272 receives information related to sub events of a touch-based gesture. When an application has multiple views organized in a hierarchy, hit view determination module 272 identifies a hit view as the lowest view in the hierarchy which should handle the sub-event. In most circumstances, the hit view is the lowest level view in which an initiating sub-event occurs (e.g., the first sub-event in the sequence of sub-events that form an event or potential event). Once the hit view is identified by the hit view determination module 272, the hit view typically receives all sub-events related to the same touch or input source for which it was identified as the hit view.

[0139] Active event recognizer determination module 273 determines which view or views within a view hierarchy should receive a particular sequence of sub-events. In some embodiments, active event recognizer determination module 273 determines that only the hit view should receive a particular sequence of sub-events. In other embodiments, active event recognizer determination module 273 determines that all views that include the physical location of a sub-event are actively involved views, and therefore determines that all actively involved views should receive a particular sequence of sub-events. In other embodiments, even if touch sub-events were entirely confined to the area associated with one particular view, views higher in the hierarchy would still remain as actively involved views.

[0140] Event dispatcher module 274 dispatches the event information to an event recognizer (e.g., event recognizer 280). In embodiments including active event recognizer determination module 273, event dispatcher module 274 delivers the event information to an event recognizer determined by active event recognizer determination module 273. In some embodiments, event dispatcher module 274 stores in an event queue the event information, which is retrieved by a respective event receiver 282.

[0141] In some embodiments, operating system 226 includes event sorter 270. Alternatively, application 236-1 includes event sorter 270. In yet other embodiments, event sorter 270 is a stand-alone module, or a part of another module stored in memory 202, such as contact/motion module 230.

[0142] In some embodiments, application 236-1 includes a plurality of event handlers 290 and one or more application views 291, each of which includes instructions for handling touch events that occur within a respective view of the application's user interface. Each application view 291 of the application 236-1 includes one or more event recognizers 280. Typically, a respective application view 291 includes a plurality of event recognizers 280. In other embodiments, one or more of event recognizers 280 are part of a separate module, such as a user interface kit (not shown) or a higher level object from which application 236-1 inherits methods and other properties. In some embodiments, a respective event handler 290 includes one or more of: data updater 276, object updater 277, GUI updater 278, and/or event data 279 received from event sorter 270. Event handler 290 utilizes or calls data updater 276, object updater 277, or GUI updater 278 to update the application internal state 292. Alternatively, one or more of the application views 291 include one or more respective event handlers 290.

Also, in some embodiments, one or more of data updater 276, object updater 277, and GUI updater 278 are included in a respective application view 291.

[0143] A respective event recognizer 280 receives event information (e.g., event data 279) from event sorter 270 and identifies an event from the event information. Event recognizer 280 includes event receiver 282 and event comparator 284. In some embodiments, event recognizer 280 also includes at least a subset of: metadata 283, and event delivery instructions 288 (which include sub-event delivery instructions).

[0144] Event receiver 282 receives event information from event sorter 270. The event information includes information about a sub-event, for example, a touch or a touch movement. Depending on the sub-event, the event information also includes additional information, such as location of the sub-event. When the sub-event concerns motion of a touch, the event information also includes speed and direction of the sub-event. In some embodiments, events include rotation of the device from one orientation to another (e.g., from a portrait orientation to a landscape orientation, or vice versa), and the event information includes corresponding information about the current orientation (also called device attitude) of the device.

[0145] Event comparator 284 compares the event information to predefined event or sub-event definitions and, based on the comparison, determines an event or sub event, or determines or updates the state of an event or sub-event. In some embodiments, event comparator 284 includes event definitions 286. Event definitions 286 contain definitions of events (e.g., predefined sequences of sub-events), for example, event 1 (287-1), event 2 (287-2), and others. In some embodiments, sub-events in an event (287) include, for example, touch begin, touch end, touch movement, touch cancellation, and multiple touching. In one example, the definition for event 1 (287-1) is a double tap on a displayed object. The double tap, for example, comprises a first touch (touch begin) on the displayed object for a predetermined phase, a first liftoff (touch end) for a predetermined phase, a second touch (touch begin) on the displayed object for a predetermined phase, and a second liftoff (touch end) for a predetermined phase. In another example, the definition for event 2 (287-2) is a dragging on a displayed object. The dragging, for example, comprises a touch (or contact) on the displayed object for a predetermined phase, a movement of the touch across touch-sensitive display 212, and liftoff of the touch (touch end). In some embodiments, the event also includes information for one or more associated event handlers 290.

[0146] In some embodiments, event definition 287 includes a definition of an event for a respective user-interface object. In some embodiments, event comparator 284 performs a hit test to determine which user-interface object is associated with a sub-event. For example, in an application view in which three user-interface objects are displayed on touch-sensitive display 212, when a touch is detected on touch-sensitive display 212, event comparator 284 performs a hit test to determine which of the three user-interface objects is associated with the touch (sub-event). If each displayed object is associated with a respective event handler 290, the event comparator uses the result of the hit test to determine which event handler 290 should

be activated. For example, event comparator **284** selects an event handler associated with the sub-event and the object triggering the hit test.

[0147] In some embodiments, the definition for a respective event (**287**) also includes delayed actions that delay delivery of the event information until after it has been determined whether the sequence of sub-events does or does not correspond to the event recognizer's event type.

[0148] When a respective event recognizer **280** determines that the series of sub-events do not match any of the events in event definitions **286**, the respective event recognizer **280** enters an event impossible, event failed, or event ended state, after which it disregards subsequent sub-events of the touch-based gesture. In this situation, other event recognizers, if any, that remain active for the hit view continue to track and process sub-events of an ongoing touch-based gesture.

[0149] In some embodiments, a respective event recognizer **280** includes metadata **283** with configurable properties, flags, and/or lists that indicate how the event delivery system should perform sub-event delivery to actively involved event recognizers. In some embodiments, metadata **283** includes configurable properties, flags, and/or lists that indicate how event recognizers interact, or are enabled to interact, with one another. In some embodiments, metadata **283** includes configurable properties, flags, and/or lists that indicate whether sub-events are delivered to varying levels in the view or programmatic hierarchy.

[0150] In some embodiments, a respective event recognizer **280** activates event handler **290** associated with an event when one or more particular sub-events of an event are recognized. In some embodiments, a respective event recognizer **280** delivers event information associated with the event to event handler **290**. Activating an event handler **290** is distinct from sending (and deferred sending) sub-events to a respective hit view. In some embodiments, event recognizer **280** throws a flag associated with the recognized event, and event handler **290** associated with the flag catches the flag and performs a predefined process.

[0151] In some embodiments, event delivery instructions **288** include sub-event delivery instructions that deliver event information about a sub-event without activating an event handler. Instead, the sub-event delivery instructions deliver event information to event handlers associated with the series of sub-events or to actively involved views. Event handlers associated with the series of sub-events or with actively involved views receive the event information and perform a predetermined process.

[0152] In some embodiments, data updater **276** creates and updates data used in application **236-1**. For example, data updater **276** updates the telephone number used in contacts module **237**, or stores a video file used in video player module. In some embodiments, object updater **277** creates and updates objects used in application **236-1**. For example, object updater **277** creates a new user-interface object or updates the position of a user-interface object. GUI updater **278** updates the GUI. For example, GUI updater **278** prepares display information and sends it to graphics module **232** for display on a touch-sensitive display.

[0153] In some embodiments, event handler(s) **290** includes or has access to data updater **276**, object updater **277**, and GUI updater **278**. In some embodiments, data updater **276**, object updater **277**, and GUI updater **278** are included in a single module of a respective application **236-1**

or application view **291**. In other embodiments, they are included in two or more software modules.

[0154] It shall be understood that the foregoing discussion regarding event handling of user touches on touch-sensitive displays also applies to other forms of user inputs to operate multifunction devices **200** with input devices, not all of which are initiated on touch screens. For example, mouse movement and mouse button presses, optionally coordinated with single or multiple keyboard presses or holds; contact movements such as taps, drags, scrolls, etc. on touchpads; pen stylus inputs; movement of the device; oral instructions; detected eye movements; biometric inputs; and/or any combination thereof are optionally utilized as inputs corresponding to sub-events which define an event to be recognized.

[0155] FIG. 3 illustrates a portable multifunction device **200** having a touch screen **212** in accordance with some embodiments. The touch screen optionally displays one or more graphics within user interface (UI) **300**. In this embodiment, as well as others described below, a user is enabled to select one or more of the graphics by making a gesture on the graphics, for example, with one or more fingers **302** (not drawn to scale in the figure) or one or more styluses **303** (not drawn to scale in the figure). In some embodiments, selection of one or more graphics occurs when the user breaks contact with the one or more graphics. In some embodiments, the gesture optionally includes one or more taps, one or more swipes (from left to right, right to left, upward and/or downward), and/or a rolling of a finger (from right to left, left to right, upward and/or downward) that has made contact with device **200**. In some implementations or circumstances, inadvertent contact with a graphic does not select the graphic. For example, a swipe gesture that sweeps over an application icon optionally does not select the corresponding application when the gesture corresponding to selection is a tap.

[0156] Device **200** also includes one or more physical buttons, such as "home" or menu button **304**. As described previously, menu button **304** is used to navigate to any application **236** in a set of applications that is executed on device **200**. Alternatively, in some embodiments, the menu button is implemented as a soft key in a GUI displayed on touch screen **212**.

[0157] In one embodiment, device **200** includes touch screen **212**, menu button **304**, push button **306** for powering the device on/off and locking the device, volume adjustment button(s) **308**, subscriber identity module (SIM) card slot **310**, headset jack **312**, and docking/charging external port **224**. Push button **306** is, optionally, used to turn the power on/off on the device by depressing the button and holding the button in the depressed state for a predefined time interval; to lock the device by depressing the button and releasing the button before the predefined time interval has elapsed; and/or to unlock the device or initiate an unlock process. In an alternative embodiment, device **200** also accepts verbal input for activation or deactivation of some functions through microphone **213**. Device **200** also, optionally, includes one or more contact intensity sensors **265** for detecting intensity of contacts on touch screen **212** and/or one or more tactile output generators **267** for generating tactile outputs for a user of device **200**.

[0158] FIG. 4 is a block diagram of an exemplary multifunction device with a display and a touch-sensitive surface in accordance with some embodiments. Device **400** need not be portable. In some embodiments, device **400** is a laptop



computer, a desktop computer, a tablet computer, a multimedia player device, a navigation device, an educational device (such as a child's learning toy), a gaming system, or a control device (e.g., a home or industrial controller). Device **400** typically includes one or more processing units (CPUs) **410**, one or more network or other communications interfaces **460**, memory **470**, and one or more communication buses **420** for interconnecting these components. Communication buses **420** optionally include circuitry (sometimes called a chipset) that interconnects and controls communications between system components. Device **400** includes input/output (I/O) interface **430** comprising display **440**, which is typically a touch screen display. I/O interface **430** also optionally includes a keyboard and/or mouse (or other pointing device) **450** and touchpad **455**, tactile output generator **457** for generating tactile outputs on device **400** (e.g., similar to tactile output generator(s) **267** described above with reference to FIG. 2A), sensors **459** (e.g., optical, acceleration, proximity, touch-sensitive, and/or contact intensity sensors similar to contact intensity sensor(s) **265** described above with reference to FIG. 2A). Memory **470** includes high-speed random access memory, such as DRAM, SRAM, DDR RAM, or other random access solid state memory devices; and optionally includes non-volatile memory, such as one or more magnetic disk storage devices, optical disk storage devices, flash memory devices, or other non-volatile solid state storage devices. Memory **470** optionally includes one or more storage devices remotely located from CPU(s) **410**. In some embodiments, memory **470** stores programs, modules, and data structures analogous to the programs, modules, and data structures stored in memory **202** of portable multifunction device **200** (FIG. 2A), or a subset thereof. Furthermore, memory **470** optionally stores additional programs, modules, and data structures not present in memory **202** of portable multifunction device **200**. For example, memory **470** of device **400** optionally stores drawing module **480**, presentation module **482**, word processing module **484**, website creation module **486**, disk authoring module **488**, and/or spreadsheet module **490**, while memory **202** of portable multifunction device **200** (FIG. 2A) optionally does not store these modules.

[0159] Each of the above-identified elements in FIG. 4 is, in some examples, stored in one or more of the previously mentioned memory devices. Each of the above-identified modules corresponds to a set of instructions for performing a function described above. The above-identified modules or programs (e.g., sets of instructions) need not be implemented as separate software programs, procedures, or modules, and thus various subsets of these modules are combined or otherwise rearranged in various embodiments. In some embodiments, memory **470** stores a subset of the modules and data structures identified above. Furthermore, memory **470** stores additional modules and data structures not described above.

[0160] Attention is now directed towards embodiments of user interfaces that can be implemented on, for example, portable multifunction device **200**.

[0161] FIG. 5A illustrates an exemplary user interface for a menu of applications on portable multifunction device **200** in accordance with some embodiments. Similar user interfaces are implemented on device **400**. In some embodiments, user interface **500** includes the following elements, or a subset or superset thereof:

[0162] Signal strength indicator(s) **502** for wireless communication(s), such as cellular and Wi-Fi signals;

[0163] Time **504**;

[0164] Bluetooth indicator **505**;

[0165] Battery status indicator **506**;

[0166] Tray **508** with icons for frequently used applications, such as:

[0167] Icon **516** for telephone module **238**, labeled "Phone," which optionally includes an indicator **514** of the number of missed calls or voicemail messages;

[0168] Icon **518** for e-mail client module **240**, labeled "Mail," which optionally includes an indicator **510** of the number of unread e-mails;

[0169] Icon **520** for browser module **247**, labeled "Browser;" and

[0170] Icon **522** for video and music player module **252**, also referred to as iPod (trademark of Apple Inc.) module **252**, labeled "iPod;" and

[0171] Icons for other applications, such as:

[0172] Icon **524** for IM module **241**, labeled "Messages;"

[0173] Icon **526** for calendar module **248**, labeled "Calendar;"

[0174] Icon **528** for image management module **244**, labeled "Photos;"

[0175] Icon **530** for camera module **243**, labeled "Camera;"

[0176] Icon **532** for online video module **255**, labeled "Online Video;"

[0177] Icon **534** for stocks widget **249-2**, labeled "Stocks;"

[0178] Icon **536** for map module **254**, labeled "Maps;"

[0179] Icon **538** for weather widget **249-1**, labeled "Weather;"

[0180] Icon **540** for alarm clock widget **249-4**, labeled "Clock;"

[0181] Icon **542** for workout support module **242**, labeled "Workout Support;"

[0182] Icon **544** for notes module **253**, labeled "Notes;" and

[0183] Icon **546** for a settings application or module, labeled "Settings," which provides access to settings for device **200** and its various applications **236**.

[0184] It should be noted that the icon labels illustrated in FIG. 5A are merely exemplary. For example, icon **522** for video and music player module **252** is optionally labeled "Music" or "Music Player." Other labels are, optionally, used for various application icons. In some embodiments, a label for a respective application icon includes a name of an application corresponding to the respective application icon. In some embodiments, a label for a particular application icon is distinct from a name of an application corresponding to the particular application icon.

[0185] FIG. 5B illustrates an exemplary user interface on a device (e.g., device **400**, FIG. 4) with a touch-sensitive surface **551** (e.g., a tablet or touchpad **455**, FIG. 4) that is separate from the display **550** (e.g., touch screen display **212**). Device **400** also, optionally, includes one or more contact intensity sensors (e.g., one or more of sensors **457**) for detecting intensity of contacts on touch-sensitive surface **551** and/or one or more tactile output generators **459** for generating tactile outputs for a user of device **400**.

[0186] Although some of the examples which follow will be given with reference to inputs on touch screen display 212 (where the touch-sensitive surface and the display are combined), in some embodiments, the device detects inputs on a touch-sensitive surface that is separate from the display, as shown in FIG. 5B. In some embodiments, the touch-sensitive surface (e.g., 551 in FIG. 5B) has a primary axis (e.g., 552 in FIG. 5B) that corresponds to a primary axis (e.g., 553 in FIG. 5B) on the display (e.g., 550). In accordance with these embodiments, the device detects contacts (e.g., 560 and 562 in FIG. 5B) with the touch-sensitive surface 551 at locations that correspond to respective locations on the display (e.g., in FIG. 5B, 560 corresponds to 568 and 562 corresponds to 570). In this way, user inputs (e.g., contacts 560 and 562, and movements thereof) detected by the device on the touch-sensitive surface (e.g., 551 in FIG. 5B) are used by the device to manipulate the user interface on the display (e.g., 550 in FIG. 5B) of the multifunction device when the touch-sensitive surface is separate from the display. It should be understood that similar methods are, optionally, used for other user interfaces described herein.

[0187] Additionally, while the following examples are given primarily with reference to finger inputs (e.g., finger contacts, finger tap gestures, finger swipe gestures), it should be understood that, in some embodiments, one or more of the finger inputs are replaced with input from another input device (e.g., a mouse-based input or stylus input). For example, a swipe gesture is, optionally, replaced with a mouse click (e.g., instead of a contact) followed by movement of the cursor along the path of the swipe (e.g., instead of movement of the contact). As another example, a tap gesture is, optionally, replaced with a mouse click while the cursor is located over the location of the tap gesture (e.g., instead of detection of the contact followed by ceasing to detect the contact). Similarly, when multiple user inputs are simultaneously detected, it should be understood that multiple computer mice are, optionally, used simultaneously, or a mouse and finger contacts are, optionally, used simultaneously.

[0188] FIG. 6A illustrates exemplary personal electronic device 600. Device 600 includes body 602. In some embodiments, device 600 includes some or all of the features described with respect to devices 200 and 400 (e.g., FIGS. 2A-4). In some embodiments, device 600 has touch-sensitive display screen 604, hereafter touch screen 604. Alternatively, or in addition to touch screen 604, device 600 has a display and a touch-sensitive surface. As with devices 200 and 400, in some embodiments, touch screen 604 (or the touch-sensitive surface) has one or more intensity sensors for detecting intensity of contacts (e.g., touches) being applied. The one or more intensity sensors of touch screen 604 (or the touch-sensitive surface) provide output data that represents the intensity of touches. The user interface of device 600 responds to touches based on their intensity, meaning that touches of different intensities can invoke different user interface operations on device 600.

[0189] Techniques for detecting and processing touch intensity are found, for example, in related applications: International Patent Application Serial No. PCT/US2013/040061, titled “Device, Method, and Graphical User Interface for Displaying User Interface Objects Corresponding to an Application,” filed May 8, 2013, and International Patent Application Serial No. PCT/US2013/069483, titled “Device, Method, and Graphical User Interface for Transi-

tioning Between Touch Input to Display Output Relationships,” filed Nov. 11, 2013, each of which is hereby incorporated by reference in their entirety.

[0190] In some embodiments, device 600 has one or more input mechanisms 606 and 608. Input mechanisms 606 and 608, if included, are physical. Examples of physical input mechanisms include push buttons and rotatable mechanisms. In some embodiments, device 600 has one or more attachment mechanisms. Such attachment mechanisms, if included, can permit attachment of device 600 with, for example, hats, eyewear, earrings, necklaces, shirts, jackets, bracelets, watch straps, chains, trousers, belts, shoes, purses, backpacks, and so forth. These attachment mechanisms permit device 600 to be worn by a user.

[0191] FIG. 6B depicts exemplary personal electronic device 600. In some embodiments, device 600 includes some or all of the components described with respect to FIGS. 2A, 2B, and 4. Device 600 has bus 612 that operatively couples I/O section 614 with one or more computer processors 616 and memory 618. I/O section 614 is connected to display 604, which can have touch-sensitive component 622 and, optionally, touch-intensity sensitive component 624. In addition, I/O section 614 is connected with communication unit 630 for receiving application and operating system data, using Wi-Fi, Bluetooth, near field communication (NFC), cellular, and/or other wireless communication techniques. Device 600 includes input mechanisms 606 and/or 608. Input mechanism 606 is a rotatable input device or a depressible and rotatable input device, for example. Input mechanism 608 is a button, in some examples.

[0192] Input mechanism 608 is a microphone, in some examples. Personal electronic device 600 includes, for example, various sensors, such as GPS sensor 632, accelerometer 634, directional sensor 640 (e.g., compass), gyroscope 636, motion sensor 638, and/or a combination thereof, all of which are operatively connected to I/O section 614.

[0193] Memory 618 of personal electronic device 600 is a non-transitory computer-readable storage medium, for storing computer-executable instructions, which, when executed by one or more computer processors 616, for example, cause the computer processors to perform the techniques and processes described below. The computer-executable instructions, for example, are also stored and/or transported within any non-transitory computer-readable storage medium for use by or in connection with an instruction execution system, apparatus, or device, such as a computer-based system, processor-containing system, or other system that can fetch the instructions from the instruction execution system, apparatus, or device and execute the instructions. Personal electronic device 600 is not limited to the components and configuration of FIG. 6B, but can include other or additional components in multiple configurations.

[0194] As used here, the term “affordance” refers to a user-interactive graphical user interface object that is, for example, displayed on the display screen of devices 200, 400, and/or 600 (FIGS. 2A, 4, and 6A-B). For example, an image (e.g., icon), a button, and text (e.g., hyperlink) each constitutes an affordance.

[0195] As used herein, the term “focus selector” refers to an input element that indicates a current part of a user interface with which a user is interacting. In some implementations that include a cursor or other location marker, the cursor acts as a “focus selector” so that when an input (e.g.,

a press input) is detected on a touch-sensitive surface (e.g., touchpad **455** in FIG. 4 or touch-sensitive surface **551** in FIG. 5B) while the cursor is over a particular user interface element (e.g., a button, window, slider or other user interface element), the particular user interface element is adjusted in accordance with the detected input. In some implementations that include a touch screen display (e.g., touch-sensitive display system **212** in FIG. 2A or touch screen **212** in FIG. 5A) that enables direct interaction with user interface elements on the touch screen display, a detected contact on the touch screen acts as a “focus selector” so that when an input (e.g., a press input by the contact) is detected on the touch screen display at a location of a particular user interface element (e.g., a button, window, slider, or other user interface element), the particular user interface element is adjusted in accordance with the detected input. In some implementations, focus is moved from one region of a user interface to another region of the user interface without corresponding movement of a cursor or movement of a contact on a touch screen display (e.g., by using a tab key or arrow keys to move focus from one button to another button); in these implementations, the focus selector moves in accordance with movement of focus between different regions of the user interface. Without regard to the specific form taken by the focus selector, the focus selector is generally the user interface element (or contact on a touch screen display) that is controlled by the user so as to communicate the user’s intended interaction with the user interface (e.g., by indicating, to the device, the element of the user interface with which the user is intending to interact). For example, the location of a focus selector (e.g., a cursor, a contact, or a selection box) over a respective button while a press input is detected on the touch-sensitive surface (e.g., a touchpad or touch screen) will indicate that the user is intending to activate the respective button (as opposed to other user interface elements shown on a display of the device).

**[0196]** As used in the specification and claims, the term “characteristic intensity” of a contact refers to a characteristic of the contact based on one or more intensities of the contact. In some embodiments, the characteristic intensity is based on multiple intensity samples. The characteristic intensity is, optionally, based on a predefined number of intensity samples, or a set of intensity samples collected during a predetermined time period (e.g., 0.05, 0.1, 0.2, 0.5, 1, 2, 5, 10 seconds) relative to a predefined event (e.g., after detecting the contact, prior to detecting liftoff of the contact, before or after detecting a start of movement of the contact, prior to detecting an end of the contact, before or after detecting an increase in intensity of the contact, and/or before or after detecting a decrease in intensity of the contact). A characteristic intensity of a contact is, optionally based on one or more of: a maximum value of the intensities of the contact, a mean value of the intensities of the contact, an average value of the intensities of the contact, a top 10 percentile value of the intensities of the contact, a value at the half maximum of the intensities of the contact, a value at the 90 percent maximum of the intensities of the contact, or the like. In some embodiments, the duration of the contact is used in determining the characteristic intensity (e.g., when the characteristic intensity is an average of the intensity of the contact over time). In some embodiments, the characteristic intensity is compared to a set of one or more intensity thresholds to determine whether an operation has been

performed by a user. For example, the set of one or more intensity thresholds includes a first intensity threshold and a second intensity threshold. In this example, a contact with a characteristic intensity that does not exceed the first threshold results in a first operation, a contact with a characteristic intensity that exceeds the first intensity threshold and does not exceed the second intensity threshold results in a second operation, and a contact with a characteristic intensity that exceeds the second threshold results in a third operation. In some embodiments, a comparison between the characteristic intensity and one or more thresholds is used to determine whether or not to perform one or more operations (e.g., whether to perform a respective operation or forgo performing the respective operation) rather than being used to determine whether to perform a first operation or a second operation.

**[0197]** In some embodiments, a portion of a gesture is identified for purposes of determining a characteristic intensity. For example, a touch-sensitive surface receives a continuous swipe contact transitioning from a start location and reaching an end location, at which point the intensity of the contact increases. In this example, the characteristic intensity of the contact at the end location is based on only a portion of the continuous swipe contact, and not the entire swipe contact (e.g., only the portion of the swipe contact at the end location). In some embodiments, a smoothing algorithm is applied to the intensities of the swipe contact prior to determining the characteristic intensity of the contact. For example, the smoothing algorithm optionally includes one or more of: an unweighted sliding-average smoothing algorithm, a triangular smoothing algorithm, a median filter smoothing algorithm, and/or an exponential smoothing algorithm. In some circumstances, these smoothing algorithms eliminate narrow spikes or dips in the intensities of the swipe contact for purposes of determining a characteristic intensity.

**[0198]** The intensity of a contact on the touch-sensitive surface is characterized relative to one or more intensity thresholds, such as a contact-detection intensity threshold, a light press intensity threshold, a deep press intensity threshold, and/or one or more other intensity thresholds. In some embodiments, the light press intensity threshold corresponds to an intensity at which the device will perform operations typically associated with clicking a button of a physical mouse or a trackpad. In some embodiments, the deep press intensity threshold corresponds to an intensity at which the device will perform operations that are different from operations typically associated with clicking a button of a physical mouse or a trackpad. In some embodiments, when a contact is detected with a characteristic intensity below the light press intensity threshold (e.g., and above a nominal contact-detection intensity threshold below which the contact is no longer detected), the device will move a focus selector in accordance with movement of the contact on the touch-sensitive surface without performing an operation associated with the light press intensity threshold or the deep press intensity threshold. Generally, unless otherwise stated, these intensity thresholds are consistent between different sets of user interface figures.

**[0199]** An increase of characteristic intensity of the contact from an intensity below the light press intensity threshold to an intensity between the light press intensity threshold and the deep press intensity threshold is sometimes referred to as a “light press” input. An increase of characteristic

intensity of the contact from an intensity below the deep press intensity threshold to an intensity above the deep press intensity threshold is sometimes referred to as a “deep press” input. An increase of characteristic intensity of the contact from an intensity below the contact-detection intensity threshold to an intensity between the contact-detection intensity threshold and the light press intensity threshold is sometimes referred to as detecting the contact on the touch-surface. A decrease of characteristic intensity of the contact from an intensity above the contact-detection intensity threshold to an intensity below the contact-detection intensity threshold is sometimes referred to as detecting liftoff of the contact from the touch-surface. In some embodiments, the contact-detection intensity threshold is zero. In some embodiments, the contact-detection intensity threshold is greater than zero.

**[0200]** In some embodiments described herein, one or more operations are performed in response to detecting a gesture that includes a respective press input or in response to detecting the respective press input performed with a respective contact (or a plurality of contacts), where the respective press input is detected based at least in part on detecting an increase in intensity of the contact (or plurality of contacts) above a press-input intensity threshold. In some embodiments, the respective operation is performed in response to detecting the increase in intensity of the respective contact above the press-input intensity threshold (e.g., a “down stroke” of the respective press input). In some embodiments, the press input includes an increase in intensity of the respective contact above the press-input intensity threshold and a subsequent decrease in intensity of the contact below the press-input intensity threshold, and the respective operation is performed in response to detecting the subsequent decrease in intensity of the respective contact below the press-input threshold (e.g., an “up stroke” of the respective press input).

**[0201]** In some embodiments, the device employs intensity hysteresis to avoid accidental inputs sometimes termed “jitter,” where the device defines or selects a hysteresis intensity threshold with a predefined relationship to the press-input intensity threshold (e.g., the hysteresis intensity threshold is X intensity units lower than the press-input intensity threshold or the hysteresis intensity threshold is 75%, 90%, or some reasonable proportion of the press-input intensity threshold). Thus, in some embodiments, the press input includes an increase in intensity of the respective contact above the press-input intensity threshold and a subsequent decrease in intensity of the contact below the hysteresis intensity threshold that corresponds to the press-input intensity threshold, and the respective operation is performed in response to detecting the subsequent decrease in intensity of the respective contact below the hysteresis intensity threshold (e.g., an “up stroke” of the respective press input). Similarly, in some embodiments, the press input is detected only when the device detects an increase in intensity of the contact from an intensity at or below the hysteresis intensity threshold to an intensity at or above the press-input intensity threshold and, optionally, a subsequent decrease in intensity of the contact to an intensity at or below the hysteresis intensity, and the respective operation is performed in response to detecting the press input (e.g., the increase in intensity of the contact or the decrease in intensity of the contact, depending on the circumstances).

**[0202]** For ease of explanation, the descriptions of operations performed in response to a press input associated with a press-input intensity threshold or in response to a gesture including the press input are, optionally, triggered in response to detecting either: an increase in intensity of a contact above the press-input intensity threshold, an increase in intensity of a contact from an intensity below the hysteresis intensity threshold to an intensity above the press-input intensity threshold, a decrease in intensity of the contact below the press-input intensity threshold, and/or a decrease in intensity of the contact below the hysteresis intensity threshold corresponding to the press-input intensity threshold. Additionally, in examples where an operation is described as being performed in response to detecting a decrease in intensity of a contact below the press-input intensity threshold, the operation is, optionally, performed in response to detecting a decrease in intensity of the contact below a hysteresis intensity threshold corresponding to, and lower than, the press-input intensity threshold.

### 3. Digital Assistant System

**[0203]** FIG. 7A illustrates a block diagram of digital assistant system 700 in accordance with various examples. In some examples, digital assistant system 700 is implemented on a standalone computer system. In some examples, digital assistant system 700 is distributed across multiple computers. In some examples, 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 one or more user devices (e.g., devices 104, 122, 200, 400, or 600) and communicates with the server portion (e.g., server system 108) through one or more networks, e.g., as shown in FIG. 1. In some examples, digital assistant system 700 is an implementation of server system 108 (and/or DA server 106) shown in FIG. 1. It should be noted that digital assistant system 700 is only one example of a digital assistant system, and that digital assistant system 700 can have more or fewer components than shown, can combine two or more components, or can have a different configuration or arrangement of the components. The various components shown in FIG. 7A are 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 thereof.

**[0204]** Digital assistant system 700 includes memory 702, one or more processors 704, input/output (I/O) interface 706, and network communications interface 708. These components can communicate with one another over one or more communication buses or signal lines 710.

**[0205]** In some examples, memory 702 includes a non-transitory computer-readable medium, such as high-speed random access memory and/or a non-volatile computer-readable storage medium (e.g., one or more magnetic disk storage devices, flash memory devices, or other non-volatile solid-state memory devices).

**[0206]** In some examples, I/O interface 706 couples input/output devices 716 of digital assistant system 700, such as displays, keyboards, touch screens, and microphones, to user interface module 722. I/O interface 706, in conjunction with user interface module 722, receives user inputs (e.g., voice input, keyboard inputs, touch inputs, etc.) and processes them accordingly. In some examples, e.g., when the digital assistant is implemented on a standalone user device, digital

assistant system **700** includes any of the components and I/O communication interfaces described with respect to devices **200**, **400**, or **600** in FIGS. 2A, 4, 6A-B, respectively. In some examples, digital assistant system **700** represents the server portion of a digital assistant implementation, and can interact with the user through a client-side portion residing on a user device (e.g., devices **104**, **200**, **400**, or **600**).

[0207] In some examples, the network communications interface **708** includes wired communication port(s) **712** and/or wireless transmission and reception circuitry **714**. The wired communication port(s) receives and send communication signals via one or more wired interfaces, e.g., Ethernet, Universal Serial Bus (USB), FIREWIRE, etc. The wireless circuitry **714** receives and sends RF signals and/or optical signals from/to communications networks and other communications devices. The wireless communications 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. Network communications interface **708** enables communication between digital assistant system **700** 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.

[0208] In some examples, memory **702**, or the computer-readable storage media of memory **702**, stores programs, modules, instructions, and data structures including all or a subset of: operating system **718**, communications module **720**, user interface module **722**, one or more applications **724**, and digital assistant module **726**. In particular, memory **702**, or the computer-readable storage media of memory **702**, stores instructions for performing the processes described below. One or more processors **704** execute these programs, modules, and instructions, and reads/writes from/to the data structures.

[0209] Operating system **718** (e.g., Darwin, RTXC, LINUX, UNIX, iOS, 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.

[0210] Communications module **720** facilitates communications between digital assistant system **700** with other devices over network communications interface **708**. For example, communications module **720** communicates with RF circuitry **208** of electronic devices such as devices **200**, **400**, and **600** shown in FIG. 2A, 4, 6A-B, respectively. Communications module **720** also includes various components for handling data received by wireless circuitry **714** and/or wired communications port **712**.

[0211] User interface module **722** receives commands and/or inputs from a user via I/O interface **706** (e.g., from a keyboard, touch screen, pointing device, controller, and/or microphone), and generate user interface objects on a display. User interface module **722** also prepares and delivers outputs (e.g., speech, sound, animation, text, icons, vibrations, haptic feedback, light, etc.) to the user via the I/O interface **706** (e.g., through displays, audio channels, speakers, touch-pads, etc.).

[0212] Applications **724** include programs and/or modules that are configured to be executed by one or more processors

**704**. For example, if the digital assistant system is implemented on a standalone user device, applications **724** include user applications, such as games, a calendar application, a navigation application, or an email application. If digital assistant system **700** is implemented on a server, applications **724** include resource management applications, diagnostic applications, or scheduling applications, for example.

[0213] Memory **702** also stores digital assistant module **726** (or the server portion of a digital assistant). In some examples, digital assistant module **726** includes the following sub-modules, or a subset or superset thereof: input/output processing module **728**, speech-to-text (STT) processing module **730**, natural language processing module **732**, dialogue flow processing module **734**, task flow processing module **736**, service processing module **738**, and speech synthesis module **740**. Each of these modules has access to one or more of the following systems or data and models of the digital assistant module **726**, or a subset or superset thereof: ontology **760**, vocabulary index **744**, user data **748**, task flow models **754**, service models **756**, and ASR systems **758**.

[0214] In some examples, using the processing modules, data, and models implemented in digital assistant module **726**, the digital assistant can perform at least some of the following: converting speech input into text; 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, games, intentions, etc.); determining the task flow for fulfilling the inferred intent; and executing the task flow to fulfill the inferred intent.

[0215] In some examples, as shown in FIG. 7B, I/O processing module **728** interacts with the user through I/O devices **716** in FIG. 7A or with a user device (e.g., devices **104**, **200**, **400**, or **600**) through network communications interface **708** in FIG. 7A to obtain user input (e.g., a speech input) and to provide responses (e.g., as speech outputs) to the user input. I/O processing module **728** optionally obtains contextual information associated with the user input from the user device, along with or shortly after the receipt of the user input. The contextual information includes user-specific data, vocabulary, and/or preferences relevant to the user input. In some examples, the contextual information also includes software and hardware states of the user device 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 examples, I/O processing module **728** also sends follow-up questions to, and receive answers from, the user regarding the user request. When a user request is received by I/O processing module **728** and the user request includes speech input, I/O processing module **728** forwards the speech input to STT processing module **730** (or speech recognizer) for speech-to-text conversions.

[0216] STT processing module **730** includes one or more ASR systems **758**. The one or more ASR systems **758** can process the speech input that is received through I/O processing module **728** to produce a recognition result. Each ASR system **758** includes a front-end speech pre-processor. The front-end speech pre-processor extracts representative features from the speech input. For example, the front-end speech pre-processor performs a Fourier transform on the speech input to extract spectral features that characterize the

speech input as a sequence of representative multi-dimensional vectors. Further, each ASR system 758 includes one or more speech recognition models (e.g., acoustic models and/or language models) and implements one or more speech recognition engines. Examples of speech recognition models include Hidden Markov Models, Gaussian-Mixture Models, Deep Neural Network Models, n-gram language models, and other statistical models. Examples of speech recognition engines include the dynamic time warping based engines and weighted finite-state transducers (WFST) based engines. The one or more speech recognition models and the one or more speech recognition engines are used to process the extracted representative features of the front-end speech pre-processor to produce intermediate recognitions results (e.g., phonemes, phonemic strings, and sub-words), and ultimately, text recognition results (e.g., words, word strings, or sequence of tokens). In some examples, the speech input is processed at least partially by a third-party service or on the user's device (e.g., device 104, 200, 400, or 600) to produce the recognition result. Once STT processing module 730 produces recognition results containing a text string (e.g., words, or sequence of words, or sequence of tokens), the recognition result is passed to natural language processing module 732 for intent deduction. In some examples, STT processing module 730 produces multiple candidate text representations of the speech input. Each candidate text representation is a sequence of words or tokens corresponding to the speech input. In some examples, each candidate text representation is associated with a speech recognition confidence score. Based on the speech recognition confidence scores, STT processing module 730 ranks the candidate text representations and provides the n-best (e.g., n highest ranked) candidate text representation(s) to natural language processing module 732 for intent deduction, where n is a predetermined integer greater than zero. For example, in one example, only the highest ranked (n=1) candidate text representation is passed to natural language processing module 732 for intent deduction. In another example, the five highest ranked (n=5) candidate text representations are passed to natural language processing module 732 for intent deduction.

[0217] More details on the speech-to-text processing are described in U.S. Utility application Ser. No. 13/236,942 for "Consolidating Speech Recognition Results," filed on Sep. 20, 2011, the entire disclosure of which is incorporated herein by reference.

[0218] In some examples, STT processing module 730 includes and/or accesses a vocabulary of recognizable words via phonetic alphabet conversion module 731. Each vocabulary word is associated with one or more candidate pronunciations of the word represented in a speech recognition phonetic alphabet. In particular, the vocabulary of recognizable words includes a word that is associated with a plurality of candidate pronunciations. For example, the vocabulary includes the word "tomato" that is associated with the candidate pronunciations of /tə'meɪroʊ/ and /tə'matou/. Further, vocabulary words are associated with custom candidate pronunciations that are based on previous speech inputs from the user. Such custom candidate pronunciations are stored in STT processing module 730 and are associated with a particular user via the user's profile on the device. In some examples, the candidate pronunciations for words are determined based on the spelling of the word and one or more linguistic and/or phonetic rules. In some examples, the

candidate pronunciations are manually generated, e.g., based on known canonical pronunciations.

[0219] In some examples, the candidate pronunciations are ranked based on the commonness of the candidate pronunciation. For example, the candidate pronunciation /tə'meɪroʊ/ is ranked higher than /tə'matou/, because the former is a more commonly used pronunciation (e.g., among all users, for users in a particular geographical region, or for any other appropriate subset of users). In some examples, candidate pronunciations are ranked based on whether the candidate pronunciation is a custom candidate pronunciation associated with the user. For example, custom candidate pronunciations are ranked higher than canonical candidate pronunciations. This can be useful for recognizing proper nouns having a unique pronunciation that deviates from canonical pronunciation. In some examples, candidate pronunciations are associated with one or more speech characteristics, such as geographic origin, nationality, or ethnicity.

For example, the candidate pronunciation /tə'meɪroʊ/ is associated with the United States, whereas the candidate pronunciation /z,21/ is associated with Great Britain. Further, the rank of the candidate pronunciation is based on one or more characteristics (e.g., geographic origin, nationality, ethnicity, etc.) of the user stored in the user's profile on the device. For example, it can be determined from the user's profile that the user is associated with the United States. Based on the user being associated with the United States, the candidate pronunciation /tə'meɪroʊ/ (associated with the United States) is ranked higher than the candidate pronunciation /z,21/ (associated with Great Britain). In some examples, one of the ranked candidate pronunciations is selected as a predicted pronunciation (e.g., the most likely pronunciation).

[0220] When a speech input is received, STT processing module 730 is used to determine the phonemes corresponding to the speech input (e.g., using an acoustic model), and then attempt to determine words that match the phonemes (e.g., using a language model). For example, if STT processing module 730 first identifies the sequence of phonemes /tə'meɪroʊ/ corresponding to a portion of the speech input, it can then determine, based on vocabulary index 744, that this sequence corresponds to the word "tomato."

[0221] In some examples, STT processing module 730 uses approximate matching techniques to determine words in an utterance. Thus, for example, the STT processing module 730 determines that the sequence of phonemes /tə'meɪroʊ/ corresponds to the word "tomato," even if that particular sequence of phonemes is not one of the candidate sequence of phonemes for that word.

[0222] Natural language processing module 732 ("natural language processor") of the digital assistant takes the n-best candidate text representation(s) ("word sequence(s)" or "token sequence(s)") generated by STT processing module 730, and attempts to associate each of the candidate text representations with one or more "actionable intents" recognized by the digital assistant. In some embodiments, as described in more detail below, STT processing module 730 attempts to associate each of the candidate text representation with one or more "candidate intents" using a false trigger mitigator (FTM). The FTM provides the candidate intents to a candidate intent evaluator (CIE), which evaluates whether the candidate intents include one or more

“actionable intent.” An “actionable intent” (or “user intent”) represents a task that can be performed by the digital assistant, and can have an associated task flow implemented in task flow models 754. 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 task flow models 754, or in other words, on the number and variety of “actionable intents” that the digital assistant recognizes. The effectiveness of the digital assistant, however, also depends on the assistant’s ability to infer the correct “actionable intent(s)” from the user request expressed in natural language.

[0223] In some examples, in addition to the sequence of words or tokens obtained from STT processing module 730, natural language processing module 732 also receives contextual information associated with the user request, e.g., from I/O processing module 728. The natural language processing module 732 optionally uses the contextual information to clarify, supplement, and/or further define the information contained in the candidate text representations received from STT processing module 730. The contextual 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. As described herein, contextual information is, in some examples, dynamic, and changes with time, location, content of the dialogue, and other factors.

[0224] In some examples, the natural language processing is based on, e.g., ontology 760. Ontology 760 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 ontology 760 defines how a parameter represented by the property node pertains to the task represented by the actionable intent node.

[0225] In some examples, ontology 760 is made up of actionable intent nodes and property nodes. Within ontology 760, 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 FIG. 7C, ontology 760 includes a “restaurant reservation” node (i.e., an actionable intent node). Property nodes “restaurant,” “date/time” (for the reservation), and “party size” are each directly linked to the actionable intent node (i.e., the “restaurant reservation” node).

[0226] 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 (i.e., the actionable intent node) through the intermediate property node “restaurant.” For another example, as shown in FIG. 7C, ontology 760 also includes a “set reminder” node (i.e., another actionable

intent node). Property nodes “date/time” (for 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 ontology 760.

[0227] An actionable intent node, along with its linked concept nodes, is 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 there between) associated with the particular actionable intent. For example, ontology 760 shown in FIG. 7C includes an example of restaurant reservation domain 762 and an example of reminder domain 764 within ontology 760. 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.” Reminder domain 764 includes the actionable intent node “set reminder,” and property nodes “subject” and “date/time.” In some examples, ontology 760 is made up of many domains. Each domain shares one or more property nodes with one or more other domains. For example, the “date/time” property node is associated with many different domains (e.g., a scheduling domain, a travel reservation domain, a movie ticket domain, etc.), in addition to restaurant reservation domain 762 and reminder domain 764.

[0228] While FIG. 7C illustrates two example domains within ontology 760, other domains include, for example, “find a movie,” “initiate a phone call,” “find directions,” “schedule a meeting,” “send a message,” and “provide an answer to a question,” “read a list,” “providing navigation instructions,” “provide instructions for a task” and so on. A “send a message” domain is associated with a “send a message” actionable intent node, and further includes property nodes such as “recipient(s),” “message type,” and “message body.” The property node “recipient” is further defined, for example, by the sub-property nodes such as “recipient name” and “message address.”

[0229] In some examples, ontology 760 includes all the domains (and hence actionable intents) that the digital assistant is capable of understanding and acting upon. In some examples, ontology 760 is modified, such as by adding or removing entire domains or nodes, or by modifying relationships between the nodes within the ontology 760.

[0230] In some examples, nodes associated with multiple related actionable intents are clustered under a “super domain” in ontology 760. For example, a “travel” super-domain includes a cluster of property nodes and actionable intent nodes related to travel. The actionable intent nodes related to travel 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 “travel” super domain) have many property nodes in common. For example, the actionable intent nodes for “airline reservation,” “hotel reservation,” “car rental,” “get directions,” and “find points of interest” share one or more of the property nodes “start location,” “destination,” “departure date/time,” “arrival date/time,” and “party size.”

[0231] In some examples, each node in ontology 760 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 are the so-called “vocabulary” associated with the node. The respective set of words and/or phrases associated with each node are stored in vocabulary index 744 in association with the property or actionable intent represented by the node. For example, returning to FIG. 7B, the vocabulary associated with the node for the property of “restaurant” 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” includes words and phrases such as “call,” “phone,” “dial,” “ring,” “call this number,” “make a call to,” and so on. The vocabulary index 744 optionally includes words and phrases in different languages.

[0232] Natural language processing module 732 receives the candidate text representations (e.g., text string(s) or token sequence(s)) from STT processing module 730, and for each candidate representation, determines what nodes are implicated by the words in the candidate text representation. In some examples, if a word or phrase in the candidate text representation is found to be associated with one or more nodes in ontology 760 (via vocabulary index 744), the word or phrase “triggers” or “activates” those nodes. Based on the quantity and/or relative importance of the activated nodes, natural language processing module 732 selects one of the actionable intents as the task that the user intended the digital assistant to perform. In some examples, the domain that has the most “triggered” nodes is selected. In some examples, the domain having the highest confidence value (e.g., based on the relative importance of its various triggered nodes) is selected. In some examples, the domain is selected based on a combination of the number and the importance of the triggered nodes. In some examples, 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.

[0233] User data 748 includes user-specific information, such as user-specific vocabulary, user preferences, user address, user’s default and secondary languages, user’s contact list, and other short-term or long-term information for each user. In some examples, natural language processing module 732 uses the user-specific information to supplement the information contained in the user input to further define the user intent. For example, for a user request “invite my friends to my birthday party,” natural language processing module 732 is able to access user data 748 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.

[0234] It should be recognized that in some examples, natural language processing module 732 is implemented using one or more machine learning mechanisms (e.g., neural networks). In particular, the one or more machine learning mechanisms are configured to receive a candidate text representation and contextual information associated with the candidate text representation. Based on the candidate text representation and the associated contextual information, the one or more machine learning mechanism are configured to determine intent confidence scores over a set of candidate actionable intents. Natural language processing module 732 can select one or more candidate actionable intents from the set of candidate actionable intents based on

the determined intent confidence scores. In some examples, an ontology (e.g., ontology 760) is also used to select the one or more candidate actionable intents from the set of candidate actionable intents.

[0235] Other details of searching an ontology based on a token string is described in U.S. Utility application Ser. No. 12/341,743 for “Method and Apparatus for Searching Using An Active Ontology,” filed Dec. 22, 2008, the entire disclosure of which is incorporated herein by reference.

[0236] In some examples, once natural language processing module 732 identifies an actionable intent (or domain) based on the user request, natural language processing module 732 generates a structured query to represent the identified actionable intent. In some examples, 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 says “Make me a dinner reservation at a sushi place at 7.” In this case, natural language processing module 732 is 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 includes parameters such as {Cuisine}, {Time}, {Date}, {Party Size}, and the like. In some examples, based on the speech input and the text derived from the speech input using STT processing module 730, natural language processing module 732 generates a partial structured query for the restaurant reservation domain, where the partial structured query includes the parameters {Cuisine=“Sushi”} and {Time=“7 pm”}. However, in this example, the user’s utterance contains insufficient information to complete the structured query associated with the domain. Therefore, other necessary parameters such as {Party Size} and {Date} is not specified in the structured query based on the information currently available. In some examples, natural language processing module 732 populates some parameters of the structured query with received contextual information. For example, in some examples, if the user requested a sushi restaurant “near me,” natural language processing module 732 populates a {location} parameter in the structured query with GPS coordinates from the user device.

[0237] In some examples, natural language processing module 732 identifies multiple candidate actionable intents for each candidate text representation received from STT processing module 730. Further, in some examples, a respective structured query (partial or complete) is generated for each identified candidate actionable intent. Natural language processing module 732 determines an intent confidence score for each candidate actionable intent and ranks the candidate actionable intents based on the intent confidence scores. In some examples, natural language processing module 732 passes the generated structured query (or queries), including any completed parameters, to task flow processing module 736 (“task flow processor”). In some examples, the structured query (or queries) for the m-best (e.g., m highest ranked) candidate actionable intents are provided to task flow processing module 736, where m is a predetermined integer greater than zero. In some examples, the structured query (or queries) for the m-best candidate actionable intents are provided to task flow processing module 736 with the corresponding candidate text representation(s).



[0238] Other details of inferring a user intent based on multiple candidate actionable intents determined from multiple candidate text representations of a speech input are described in U.S. Utility application Ser. No. 14/298,725 for “System and Method for Inferring User Intent From Speech Inputs,” filed Jun. 6, 2014, the entire disclosure of which is incorporated herein by reference.

[0239] Task flow processing module 736 is configured to receive the structured query (or queries) from natural language processing module 732, complete the structured query, if necessary, and perform the actions required to “complete” the user’s ultimate request. In some examples, the various procedures necessary to complete these tasks are provided in task flow models 754. In some examples, task flow models 754 include procedures for obtaining additional information from the user and task flows for performing actions associated with the actionable intent.

[0240] As described above, in order to complete a structured query, task flow processing module 736 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, task flow processing module 736 invokes dialogue flow processing module 734 to engage in a dialogue with the user. In some examples, dialogue flow processing module 734 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 I/O processing module 728. In some examples, dialogue flow processing module 734 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 task flow processing module 736 invokes dialogue flow processing module 734 to determine the “party size” and “date” information for the structured query associated with the domain “restaurant reservation,” dialogue flow processing module 734 generates questions such as “For how many people?” and “On which day?” to pass to the user. Once answers are received from the user, dialogue flow processing module 734 then populates the structured query with the missing information, or pass the information to task flow processing module 736 to complete the missing information from the structured query.

[0241] Once task flow processing module 736 has completed the structured query for an actionable intent, task flow processing module 736 proceeds to perform the ultimate task associated with the actionable intent. Accordingly, task flow processing module 736 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” 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 Café, date=3/12/2012, time=7 pm, party size=5}, task flow processing module 736 performs the steps of: (1) logging onto a server of the ABC Café 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.

[0242] In some examples, task flow processing module 736 employs the assistance of service processing module 738 (“service processing module”) to complete a task requested in the user input or to provide an informational answer requested in the user input. For example, service processing module 738 acts on behalf of task flow processing module 736 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 examples, the protocols and application programming interfaces (API) required by each service are specified by a respective service model among service models 756. Service processing module 738 accesses the appropriate service model for a service and generate requests for the service in accordance with the protocols and APIs required by the service according to the service model.

[0243] For example, if a restaurant has enabled an online reservation service, the restaurant submits a service model 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 task flow processing module 736, service processing module 738 establishes 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.

[0244] In some examples, natural language processing module 732, dialogue flow processing module 734, and task flow processing module 736 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 (i.e., an output to the user, or the completion of a task) to fulfill the user’s intent. The generated response is a dialogue response to the speech input that at least partially fulfills the user’s intent. Further, in some examples, the generated response is output as a speech output. In these examples, the generated response is sent to speech synthesis module 740 (e.g., speech synthesizer) where it can be processed to synthesize the dialogue response in speech form. In yet other examples, the generated response is data content relevant to satisfying a user request in the speech input.

[0245] In examples where task flow processing module 736 receives multiple structured queries from natural language processing module 732, task flow processing module 736 initially processes the first structured query of the received structured queries to attempt to complete the first structured query and/or execute one or more tasks or actions represented by the first structured query. In some examples, the first structured query corresponds to the highest ranked actionable intent. In other examples, the first structured query is selected from the received structured queries based on a combination of the corresponding speech recognition confidence scores and the corresponding intent confidence scores. In some examples, if task flow processing module 736 encounters an error during processing of the first structured query (e.g., due to an inability to determine a necessary parameter), the task flow processing module 736 can proceed to select and process a second structured query of the received structured queries that corresponds to a lower

ranked actionable intent. The second structured query is selected, for example, based on the speech recognition confidence score of the corresponding candidate text representation, the intent confidence score of the corresponding candidate actionable intent, a missing necessary parameter in the first structured query, or any combination thereof.

[0246] Speech synthesis module 740 is configured to synthesize speech outputs for presentation to the user. Speech synthesis module 740 synthesizes speech outputs based on text provided by the digital assistant. For example, the generated dialogue response is in the form of a text string. Speech synthesis module 740 converts the text string to an audible speech output. Speech synthesis module 740 uses any appropriate speech synthesis technique in order to generate speech outputs from text, including, but not limited, to concatenative synthesis, unit selection synthesis, diphone synthesis, domain-specific synthesis, formant synthesis, articulatory synthesis, hidden Markov model (HMM) based synthesis, and sinewave synthesis. In some examples, speech synthesis module 740 is configured to synthesize individual words based on phonemic strings corresponding to the words. For example, a phonemic string is associated with a word in the generated dialogue response. The phonemic string is stored in metadata associated with the word. Speech synthesis model 740 is configured to directly process the phonemic string in the metadata to synthesize the word in speech form.

[0247] In some examples, instead of (or in addition to) using speech synthesis module 740, speech synthesis is performed on a remote device (e.g., the server system 108), and the synthesized speech is sent to the user device for output to the user. For example, this can occur in some implementations where outputs for a digital assistant are generated at a server system. And because server systems generally have more processing power or resources than a user device, it is possible to obtain higher quality speech outputs than would be practical with client-side synthesis.

[0248] Additional details on digital assistants can be found in the U.S. Utility application Ser. No. 12/987,982, entitled “Intelligent Automated Assistant,” filed Jan. 10, 2011, and U.S. Utility application Ser. No. 13/251,088, entitled “Generating and Processing Task Items That Represent Tasks to Perform,” filed Sep. 30, 2011, the entire disclosures of which are incorporated herein by reference.

#### 4. Exemplary Architecture and Functionality of a Virtual Assistant

[0249] FIG. 8 illustrates a block diagram of a virtual assistant 800 for providing natural language interaction. In some examples, virtual assistant 800 (e.g., digital assistant system 700) can be implemented by a user device according to various embodiments. In some embodiments, the user device, a server (e.g., server 108), or a combination thereof, can implement virtual assistant 800. The user device can be implemented using, for example, device 104, 200, 400, 900, 1120, 1220, 1320, or 1420 as illustrated in FIGS. 1, 2A-2B, 4, 9, 11A-11B, 12A-12D, 13A-13B, and 14A-14D. In some embodiments, as illustrated in FIG. 8, virtual assistant 800 includes the following sub-modules, or a subset or superset thereof: an input module 810, a natural language engine 820, a false trigger mitigator (FTM) 840, a candidate intent evaluator (CIE) 860, and a task execution module 880. In some embodiments, virtual assistant 800 can be implemented using digital assistant module 726 of digital assistant

system 700 shown in FIG. 7B. For example, virtual assistant 800 can include one or more modules, models, applications, vocabularies, and user data similar to those of digital assistant module 726. With reference to FIGS. 7B and 8, as one example, input module 810 can be, for example, a sub-module or a variation of Input/Output processing module 728. Natural language engine 820 can include, for example, one or more of STT processing module 730, phonetic alphabet conversion module 731, vocabulary 744, user data 748, and/or natural language processing module 732. Virtual assistant 800 can also include modules, models, applications, vocabularies, and user data that are not included in digital assistant module 726. For example, false trigger mitigator 840 as shown in FIG. 8 may not be included in digital assistant module 726.

[0250] With reference to FIG. 8, in some embodiments, input module 810 receives one or more audio streams 802. An audio stream can include one or more utterances. In some embodiments, an utterance in an audio stream can include a word, a phrase that includes a plurality of words, and/or one or more sentences. As an example illustrated in FIG. 9, an audio stream 912 may include one or more user utterances as sentences such as “It’s dark outside. Turn on the light, Siri.”

[0251] With reference back to FIG. 8, input module 810 can be activated or remain active (e.g., for a pre-configured period of time) to receive one or more audio streams 802. For example, a voice activity detector 814 of input module 810 can detect the presence or absence of human voice (e.g., based on the amplitude and/or the frequency spectrum of the input signals received at voice activity detector 814). In some embodiments, input module 810 can include a buffer 812 (e.g., a ring buffer) configured to store one or more received audio streams (e.g., store 10 seconds of audio streams).

[0252] In some embodiments, using the received audio streams stored in buffer 812, voice activity detector 814 can determine whether an audio stream includes a lexical trigger. A lexical trigger can include a single word or a plurality of words. For example, a lexical trigger can be “Hey Siri,” “Hey Assistant,” “Siri,” “Assistant,” or the like. In some embodiments, upon receiving a lexical trigger by input module 810 (e.g., the voice activity detector 814), at least a portion of virtual assistant 800 (e.g., NLE 820) is activated. In some embodiments, at least a portion of virtual assistant 800 (e.g., input module 810, NLE 820, FTM 840) can remain active or remain active for a pre-configured period of time. For example, input module 810 can remain active to receive audio streams 802 for 30 seconds. In the embodiments where at least a portion of virtual assistant 800 remain active for at least a pre-configured period of time, a lexical trigger may not cause a portion of virtual assistant 800 to become activated. Instead, a lexical trigger may indicate that at least a portion of an audio stream is directed to the virtual assistant.

[0253] Using the above example shown in FIG. 9, the utterance “Turn on the light, Siri” includes a single-word lexical trigger “Siri,” which indicates that the utterance is directed to the virtual assistant operating on device 900. In some embodiments, a lexical trigger can be positioned at the beginning portion of an utterance in an audio stream (e.g., “Siri, what is the stock price?”). In some embodiments, a lexical trigger can be positioned at any portion other than the beginning portion of the utterance (e.g., “Turn on the light,

Siri.”). Allowing the lexical trigger to be positioned at any portion other than the beginning portion of the utterance enables a more natural human-machine interaction between a user and the virtual assistant, rather than requiring the user to lead every utterance directed to the virtual assistant with the lexical trigger. This enhances the user experience, reduces power consumption, and improves system efficiency.

[0254] With reference back to FIG. 8, as described above, input module 810 receives one or more audio streams 802. Virtual assistant 800 (e.g., using VAD 814 of input module 810) can determine the beginning point/ending point and/or the duration of a particular audio stream of the one or more audio streams 802. Based on such determination, virtual assistant 800 can determine whether a lexical trigger is included within the particular audio stream.

[0255] As one example, to detect the beginning point of a particular audio stream, virtual assistant 800 detects, via a microphone (not shown in FIG. 8), an absence of voice activity before receiving the particular audio stream (e.g., detecting a silence or a pause between adjacent utterances). Virtual assistant 800 can further determine whether the absence of voice activity before receiving the particular audio stream exceeds a first threshold period of time (e.g., 3 seconds). If the absence of voice activity exceeds the first threshold period of time, virtual assistant 800 determines the beginning point of the particular audio stream. For instance, if the first threshold period of time is 3 seconds and if virtual assistant 800 determines that before receiving the first audio stream, there is a 5-second silence or that input module 810 does not receive any audio input for the past 5 seconds, virtual assistant 800 determines that the period of absence of voice activity exceeds the first threshold period of time. As a result, the beginning point of the particular audio stream can be determined.

[0256] In some embodiments, to detect an end point of a particular audio stream, virtual assistant 800 can detect, via a microphone (not shown in FIG. 8), an absence of voice activity after receiving one or more utterances of the particular audio stream (e.g., detecting a silence or a pause between adjacent utterances). Virtual assistant 800 can further determine whether the absence of voice activity after receiving the particular audio stream exceeds a second threshold period of time (e.g., 3 seconds). If the period of absence of voice activity exceeds the second threshold period of time, virtual assistant 800 determines the end point of the first audio stream. For instance, if the second threshold period of time is 3 seconds and if virtual assistant 800 determines that after receiving the particular audio stream, there is a 5-second silence or that input module 810 does not receive any further audio input for the next 5 seconds, virtual assistant 800 can determine that the period of absence of voice activity exceeds the second threshold period of time. As a result, the end point of the particular audio stream can be determined.

[0257] As an example shown in FIG. 9, a particular audio stream 912 may include a first utterance (e.g., “It’s dark outside.”) and a second utterance (e.g., “Turn on the light, Siri.”). The virtual assistant operating on device 900 (e.g., virtual assistant 800) can determine that before receiving the first utterance of audio stream 912, no audio input is received for at least 3 seconds; and that after receiving the second utterance audio stream 912, no audio input is received for at least 3 seconds. The virtual assistant may

further determine that between the first utterance and the second utterance, there is only a short period of absence of voice activity (e.g., a short pause of 0.5 second). As a result, the virtual assistant can determine the beginning point/end point and/or the duration of audio stream 912.

[0258] In some embodiments, with reference to FIG. 8, an end point of an audio stream can be determined based on a pre-configured duration that input module 810 is configured to receive an audio stream. For example, if one or more user utterances in an audio stream are directed to a virtual assistant (e.g., to obtain information or to instruct the virtual assistant to perform a task), the user utterances typically may not last for more than 30 seconds. Accordingly, the pre-configured duration can be set to be, for example, 30 seconds. As a result, any user utterances received within 30 seconds following the detected beginning point of the particular audio stream can be determined to be utterances of a same audio stream. Accordingly, to determine an end point of a particular audio stream, in some embodiments, virtual assistant 800 (e.g., using input module 810) can detect the beginning point of an audio stream, obtain a pre-configured duration for an audio stream (e.g., 30 seconds), and determine the end point of the audio stream based on the pre-configured duration and the beginning point (e.g., the end point is 30 seconds after the beginning point).

[0259] With reference to FIG. 8, in some embodiments, an end point of an audio stream can be determined based on the capacity of buffer 812. Buffer 812 can store the received one or more audio streams as an audio file. The storage size of the audio file buffer 812 depends on the capacity of buffer 812 (e.g., in the range of megabytes). As an example of determining the end point of an audio stream based on the capacity of buffer 812, virtual assistant 800 (e.g., using input module 810) can determine a size of an audio file representing a received audio stream including one or more utterances and compare the size of the audio file with a capacity of buffer 812. In some embodiments, if the size of the audio file reaches the capacity of buffer 812 (e.g., substantially equal to the capacity of buffer 812), virtual assistant 800 can determine that the utterances included in the audio file represent an entire audio stream. Based on this determination result, virtual assistant 800 can determine the end point of the audio stream.

[0260] In some embodiments, virtual assistant 800 can also estimate a likelihood that a particular absence of voice activity before or after receiving a particular audio stream corresponds to a beginning point or an end point of the particular audio stream, respectively. As described above, virtual assistant 800 can detect an absence of voice activity before or after receiving a speech input. In some examples, an automatic speech recognition system (e.g., system 758 shown in FIG. 7A) of virtual assistant 800 processes the speech input and produces a recognition result. Based on the recognition result (e.g., text of the speech input) and one or more language models (e.g., models used by the ASR system 758 as described above), virtual assistant 800 estimates a likelihood that a particular absence of voice activity corresponds to a beginning point or an end point of the particular audio stream. In the above example of audio stream 912 (e.g., “It’s dark outside. Turn on the light, Siri.”), based on the recognition results of an ASR system and one or more language models, virtual assistant 800 may estimate the likelihood that an absence of voice activity (e.g., 3 seconds of silence) after receiving audio stream 912 corresponds to

an end point of audio stream **912**. Based on the estimated likelihood, virtual assistant **800** can detect the end point of audio stream **912** with an increased or improved confidence (e.g., based on comparison of the estimated likelihood to a threshold).

[0261] Upon detecting the end point and the beginning point of a particular audio stream, virtual assistant **800** can determine the duration of the particular audio stream. In some embodiments, based on the duration of the particular audio stream, virtual assistant **800** can determine whether a lexical trigger is included in the particular audio stream. In some embodiments, virtual assistant **800** does not determine the duration of the particular audio stream, but can determine whether a lexical trigger is included in the particular audio stream using the detected end point and the detected beginning point of the particular audio stream.

[0262] As described above, FIG. 9 illustrates an audio stream **912** that includes one or more utterances (e.g., “It’s dark outside. Turn on the light, Siri.”) and virtual assistant **800** determines that a lexical trigger (e.g., “Siri”) is included in audio stream **912**. As shown in FIG. 9, the one or more utterances of audio stream **912** may include at least one utterance that is not directed to virtual assistant **800**. For example, the first utterance “It’s dark outside” may be a comment and is not an utterance directed to the virtual assistant operating on device **900** (e.g., virtual assistant **800**). The virtual assistant may thus not perform a task or take an action based on such an utterance. As described in detail below, using various techniques, a virtual assistant can determine that an utterance is not directed to it and thus is to be disregarded. Such utterances can include, for example, user’s comments, an utterance that is directed from one user to another, or the like. The ability for a virtual assistant to determine whether an utterance is directed to it enhances a device’s operational efficiency, because the virtual assistant can disregard any utterances that are not directed to it while eliminating or reducing the requirements for leading every utterance directed to the virtual assistant with a trigger word or phrase.

[0263] With reference to FIG. 8, in some embodiments, in accordance with a determination that a particular audio stream of one or more audio streams **802** includes a lexical trigger (e.g., “Siri”), input module **810** generates one or more speech results **816**. Speech results **816** can include audio representations (e.g., phonetic representations) of the one or more utterances included in a particular audio stream. Input module **810** further provide speech results **816** to natural language engine **820**, which can include one or more modules such as speech-to-text (STT) processing module **730** and/or natural language process module **732** (shown in FIG. 7B), or a variation thereof. Based on speech results **816**, natural language engine **820** can generate candidate text representations **822** representing the one or more utterances in the particular audio stream. As described above, each candidate text representation can be a sequence of words or tokens corresponding to the utterances in the particular audio stream. Using audio stream **912** shown in FIG. 9 as an example, a natural language engine of the virtual assistant operating on device **900** can perform speech-to-text conversion of each of the utterances of audio stream **912** and generate candidate text representations including a first candidate text representation (e.g., “It’s dark outside.”) and a second candidate text representation (e.g., “Turn on the light, Siri.”).

[0264] In some embodiments, natural language engine **820** can further determine confidence levels corresponding to the one or more candidate text representations **822**. For example, as described above, each candidate text representation can be associated with a speech recognition confidence score. And natural language engine **820** can rank (e.g., using STT processing module **730**) candidate text representations **822** and provide the n-best (e.g., n highest ranked) candidate text representation(s) for candidate intent generation or derivation.

[0265] As described above, natural language engine **820** can include, for example, one or more of STT processing module **730**, phonetic alphabet conversion module **731**, vocabulary **744**, user data **748**, and/or natural language processing module **732**. In some embodiments, using natural language processing module **732**, natural language engine **820** can further interpret the candidate text representations to derive pre-mitigation intents and optionally confidence levels associated with the pre-mitigation intents. For example, natural language engine **820** can optionally rank confidence levels associated with the pre-mitigation intents and provide the n-best pre-mitigation intents for candidate intent generation or derivation.

[0266] With reference to FIG. 8, a false trigger mitigator (FTM) **840** of virtual assistant **800** can determine whether at least one candidate text representation of the one or more candidate text representations **822** (or ranked candidate text representations) is to be disregarded by virtual assistant **800**. FTM **840** can be implemented using, for example, a submodule, or a variation, of digital assistant module **726** shown in FIG. 7B. For example, FTM **840** can include a natural language processing **732**, or a variation thereof, to derive one or more candidate intents **842**. In some embodiments, FTM **840** can include a decision tree such as a simple decision tree or a boosted decision tree. As described above, in some embodiments, at least one utterance of the one or more utterances in a particular audio stream may not be directed to virtual assistant **800** and thus can be disregarded by virtual assistant **800**.

[0267] FIG. 9 illustrates one example of determining whether at least one candidate text representation is to be disregarded. As shown in FIG. 9 and described above, based on the utterances of user **910**, a natural language engine (e.g., NLE **820**) of the virtual assistant operating on device **900** can generate, for example, two candidate text representations such as “It’s dark outside.” and “Turn on the light, Siri.” In some embodiments, an FTM (e.g., FTM **840**) of the virtual assistant determines, for each of the two candidate text representations, whether the candidate text representation includes a lexical trigger. For example, the FTM determines that the first candidate text representation (e.g., “It’s dark outside.”) does not include a lexical trigger, but the second candidate text representation (e.g., “Turn on the light, Siri.”) includes a lexical trigger (e.g., “Siri”).

[0268] With reference to FIG. 8, in some embodiments, if FTM **840** determines that a candidate text representation includes a lexical trigger, FTM **840** determines that the corresponding user utterance is directed to virtual assistant **800** and therefore the particular candidate text representation is not to be disregarded. The particular candidate text representation would thus be further processed as described in more detail below. As an example shown in FIG. 9, an FTM of the virtual assistant operating on device **900** determines that the second candidate text representation (e.g.,

“Turn on the light, Siri.”) is not to be disregarded because it includes a single-word lexical trigger “Siri.”

[0269] With reference to FIG. 8, in some embodiments, if FTM 840 determines that a candidate text representation does not include a lexical trigger, FTM 840 can estimate a likelihood that the utterance corresponding to the particular candidate text representation is not directed to virtual assistant 800. Using the example shown in FIG. 9, as described above, an FTM of the virtual assistant operating on device 900 determines that the first candidate text representation (e.g., “It’s dark outside.”) does not include a lexical trigger. Further, using a decision tree, the FTM of the virtual assistant can estimate a likelihood that the utterance corresponding to the first candidate text representation is not directed to the virtual assistant. For example, the FTM can determine that the pre-mitigation intent corresponding to the first candidate text representation (e.g., “It’s dark outside.”) is not or likely not associated with a domain or a candidate intent recognized by the virtual assistant. As a result, the FTM can estimate the likelihood that the utterance corresponding to the first candidate text representation (e.g., “It’s dark outside.”) is not directed to the virtual assistant, and determine whether the estimated likelihood satisfy a threshold condition. If so, the FTM determines that the utterance is not directed to the virtual assistant. As a result, the first candidate text representation (e.g., “It’s dark outside.”) can be disregarded for the purpose of generating a candidate intent. In some embodiments, the determination of whether a candidate text representation is to be disregarded can also be based on context information such as usage pattern and/or sensory data, as described in more detail below.

[0270] With reference back to FIG. 8, in some embodiments, FTM 840 can determine, for each of candidate text representations 822, whether a particular candidate text representation is to be disregarded. If at least one of the candidate text representations 822 is to be disregarded, FTM 840 can generate one or more candidate intents 842 based on the candidate text representations that are not to be disregarded. For example, based on the pre-mitigation intents of the candidate text representations that are not to be disregarded, FTM 840 can derive candidate intents 842. Continuing the above example shown in FIG. 9, an FTM of the virtual assistant operating on device 900 determines that the first candidate text representation (e.g., “It’s dark outside.”) is to be disregarded but the second candidate text representation (e.g., “Turn on the light, Siri.”) is not to be disregarded. Accordingly, in some embodiments, the FTM of the virtual assistant can attempt to associate the pre-mitigation intent of the second candidate text representation with one of recognized domain in an ontology (e.g., ontology 760) to interpret it to derive a candidate intent. For example, the FTM can select the pre-mitigation intent (e.g., to turn on the light in the user’s living room) corresponding to the second candidate text representation as a candidate intent, while filtering out the pre-mitigation intent corresponding to the first candidate text representation.

[0271] With reference to FIG. 8, in some embodiments, FTM 840 can obtain confidence levels corresponding to each of the one or more candidate intents 842. As described above, for each of the pre-mitigation intents, NLE 820 may generate a confidence level. For example, a particular candidate text representation may be associated with multiple recognized domains and thus multiple pre-mitigation intents may be derived from the same candidate text representation.

NLE 820 and/or FTM 840 can determine a confidence score (e.g., based on the relative importance of its various triggered nodes in the ontology) to select a domain that has the highest confidence value, and in turn, determine a corresponding candidate intent.

[0272] With reference to FIG. 8, in some embodiments, virtual assistant 800 includes a candidate intent evaluator (CIE) 860. CIE 860 can determine whether one or more candidate intents 842 include at least one actionable intent. For example, CIE 860 can determine, for each of the candidate intents 842, whether a task can be performed. In some embodiments, CIE 860 can make such determination without actually performing the task (e.g., to make a dry run). As one example, CIE 860 can include one or more submodules, or variations thereof, of digital assistant module 726, such as task flow processing module 736 and dialog processing module 734 as shown in FIG. 7B. CIE 860 can perform the determination of whether a candidate intent 842 is an actionable intent using these submodules, or variations thereof, without providing the results to the speech synthesis module 740 and without outputting the results.

[0273] As another example shown in FIG. 8, CIE 860 can include task flow processing module 862 and dialog processing module 864 for determining whether a candidate intent 842 is an actionable intent. Task flow processing module 862 and dialog processing module 864 shown in FIG. 8 can be duplicates of task flow processing module 736 and dialog processing module 734 as shown in FIG. 7B, except task flow processing module 862 and dialog processing module 864 may not provide their processing results to a speech synthesis module or otherwise output the results to the user. As a result, CIE 860 can determine whether a candidate intent 842 is an actionable intent and such determination can be made without virtual assistant 8000 actually performing the task. Thus, if a task cannot be performed or can only be partially performed, CIE 860 can make such determination without causing virtual assistant 800 to actually perform the task, but rather would attempt to obtain more information (e.g., obtaining context information or initiating a dialog with the user to obtain more information) for actually performing the task. Determining whether a candidate intent is an action intent without actually performing a task improves the operational efficiency and enhances the human-machine interface. For example, such determination saves power and avoids outputting a partially performed task or causing confusion to the user.

[0274] Continuing the example shown in FIG. 9, as described above, based on a candidate text representation (e.g., “Turn on the light, Siri.”), an FTM of the virtual assistant operating on device 900 generates a candidate intent of turning on the light in the user’s living room. A CIE (e.g., CIE 860) of the virtual assistant determines whether a task can be performed with respect to this candidate intent. In some embodiments, the CIE can perform such determination using a task flow processing module (e.g., task flow processing module 862) and/or a dialog processing module (e.g., dialog processing module 864). For instance, the task flow processing module may attempt to perform the task of turning on lights in the user’s living room based on a structured query generated using the corresponding candidate intent. In some examples, the task flow processing module may determine that such task cannot be performed because it cannot find a home automation device that controls the lights in the living room. In some examples, the

task flow processing module may determine that a home automation device that controls the lights is available and therefore a corresponding task can be performed.

[0275] With reference to FIG. 8, in some embodiments, if CIE 860 determines that a task corresponding to a candidate intent 842 can be performed, it determines that the particular candidate intent 842 is an actionable intent. CIE 860 can repeat the determination for each of candidate intents 842 and generate actionable intents 868. As illustrated in FIG. 8, CIE 860 provides each of actionable intents 868 to task execution module 880, which performs one or more tasks according to actionable intents 868. Task execution module 880 can include one or more submodules, or variations thereof, of digital assistant module 726, such as task flow processing module 736 and dialog processing module 734 as shown in FIG. 7B. Unlike CIE 860, which does not cause a task to be actually performed, task execution module 880 receives a structured query (or queries) according to actionable intents 868, completes the structured query, if necessary, and performs the tasks required to “complete” the user’s ultimate request. Continuing the example shown in FIG. 9, upon receiving an actionable intent of turning on the light in the user’s living room, a task execution module of the virtual assistant operating on device 900 can cause the light to turn on and optionally output a result of execution of the actionable intent (e.g., provide an audio and/or visual output 942 such as “Your light is on.”).

[0276] With reference back to FIG. 8, in some embodiments, CIE 860 can determine whether a task can be performed based on an estimation of a confidence level associated with performing the task. For example, CIE 860 can estimate a confidence level associated with performing a task and determine whether the confidence level associated with performing the task satisfies a threshold confidence level. If CIE 860 determines that the confidence level associated with performing the task satisfies the threshold confidence level, CIE 860 determines that the task can be performed.

[0277] FIG. 10 illustrates a block diagram of an exemplary virtual assistant 1000 for providing natural language interaction using execution results and context information. In some examples, virtual assistant 1000 (e.g., digital assistant system 700) can be implemented by a user device according to various embodiments. In some embodiments, the user device, a server (e.g., server 108), or a combination thereof, can implement virtual assistant 1000. The user device can be implemented using, for example, device 104, 200, 400, 900, 1120, 1220, 1320, or 1420 as illustrated in FIGS. 1, 2A-2B, 4, 9, 11A-11B, 12A-12D, 13A-13B, and 14A-14D. Similar to virtual assistant 800, virtual assistant 1000 includes input module 810, natural language engine 816, FTM 840, CIE 860, and task execution module 880. These modules or components of virtual assistant 1000 are similar to those described above with respect to virtual assistant 800, and are therefore not repeatedly described. In some embodiments, one or more of the various modules, models, applications, vocabularies, and user data of virtual assistant 1000 can receive additional information, such as previous execution results and/or context information, to assist in voice activity detection, candidate text representations determination, candidate intent determination, actionable intent determination, and/or execution of the actionable intents.

[0278] With reference to FIGS. 10 and 11A, for example, a virtual assistant (e.g., virtual assistant 1000) operating on

device 1120 receives, via an input module (e.g., input module 810), an audio stream 1112 from user 1110. Audio stream 1112 includes one or more user utterances such as “How is the weather today, Siri? And do you know what the stock price is?” The virtual assistant generates speech results using the utterances included in audio stream 1112. Based on the speech results, the virtual assistant generates, via a natural language engine (e.g., natural language engine 820), one or more candidate text representations (e.g., candidate text representations 1022).

[0279] Next, for each candidate text representation of the candidate text representations, the virtual assistant determines, via an FTM (e.g., FTM 840), whether the particular candidate text representation is to be disregarded by the virtual assistant. As described above, in some embodiments, the virtual assistant can determine, for example, whether a particular candidate text representation includes a lexical trigger. As an example shown in FIG. 11A, the candidate text representations representing audio stream 1112 include a first candidate text representation (e.g., “How is the weather today, Siri?”) and a second candidate text representation (e.g., “And do you know what is the stock price?”). The virtual assistant determines that the first candidate text representation includes a lexical trigger (e.g., “Siri”) and therefore is not to be disregarded because the corresponding utterance is directed to the virtual assistant operating on device 1120. With respect to the second candidate text representation, the virtual assistant determines that it does not include a lexical trigger. In accordance with such a determination, the virtual assistant estimates a likelihood that the utterance corresponding to the second candidate text representation is not directed to virtual assistant 1000.

[0280] With reference to FIGS. 10 and 11A, in some embodiments, to estimate the likelihood that the utterance corresponding to the second candidate text representation is not directed to the virtual assistant, the virtual assistant obtains context information (e.g., execution results and context information 1044). The context information can be associated with a usage pattern of the virtual assistant operating on device 1120. A usage pattern of the virtual assistant can indicate a pattern of a particular activity that user 1110 performs using the virtual assistant and/or using device 1120 as shown in FIG. 11A. For example, user 1110 may frequently ask the same question about weather and stock price around 6 AM in the morning. Therefore, a usage pattern may be generated with respect to the user’s activity (e.g., asking virtual assistant of the same question) at or around 6 AM in the morning. The context information can include such a usage pattern and a time 1122 (e.g., indicating it is about 6 AM in the morning) obtained from device 1120. Based on the context information associated with the usage pattern, the virtual assistant, via an FTM (e.g., FTM 840) can estimate the likelihood that the utterance corresponding to the second candidate text representation is not directed to the virtual assistant. For example, based on the context information indicating that the user frequently asks the virtual assistant about the stock price around 6 AM in the morning and that the current time is about 6 AM, the virtual assistant can estimate that the likelihood the utterance “And do you know what is the stock price?” is not directed to the virtual assistant 1000 is low (e.g., does not satisfy a threshold). As a result, the virtual assistant determines, based on the estimated likelihood, that the second candidate text representation is not to be disregarded.

[0281] In some embodiments, context information can also be used for determining whether a candidate intent is an actionable intent. Continuing the above example shown in FIG. 11A, based on the determination that both the first candidate text representation (e.g., “How is the weather today, Siri?”) and the second candidate text representation (e.g., “And do you know what is the stock price?”) are not to be disregarded, the virtual assistant operating on device 1120 generates one or more candidate intents. For example, the virtual assistant generates a first candidate intent of obtaining weather information for today and a second candidate intent of obtaining stock price information. For each of the candidate intents, the virtual assistant, via a CIE (e.g., CIE 860), determines whether a task can be performed. As described above, a CIE can make such determination by using, for example, a task flow processing module and/or a dialog processing module without actually performing the task.

[0282] In some embodiments, with respect to a first candidate intent of obtaining weather information, the virtual assistant determines that additional information of location may be necessary to perform a task of obtaining weather information. This location information, however, is not presented in or provided by the first candidate text representation. Accordingly, in some embodiments, the virtual assistant can obtain context information associated with sensory data from one or more sensors communicatively coupled to device 1120 and determine whether the task can be performed based on the context information. For example, location data may be obtained as the context information indicating the current location of device 1120. With this context information, the virtual assistant determines that the first candidate intent of obtaining weather information is an actionable intent, because weather information can be obtained from an internal or external data source (e.g., from a weather information website) for the current location of device 1120.

[0283] With reference to FIG. 11A, with respect to a second candidate intent of obtaining stock price information, the virtual assistant determines that additional information of the name of the stock may be necessary to perform a task of obtaining stock price. This stock name information, however, is not presented in or provided by the second candidate text representation. Accordingly, in some embodiments, the virtual assistant can obtain context information associated with a usage pattern of the virtual assistant and/or device 1120. A usage pattern of the virtual assistant and/or device 1120 can indicate a pattern of a particular activity that user 1110 performs using virtual assistant 1000 and/or using the user’s device 1120 as shown in FIG. 11A. For example, user 1110 may frequently ask the question about the S&P index around 6 AM in the morning. Therefore, a usage pattern may be generated with respect to the user’s activity (e.g., asking the virtual assistant of the same question of S&P 500 index) at or around 6 AM in the morning. The context information can include such a usage pattern and a time 1122 (e.g., indicating it is about 6 AM in the morning) provided by device 1120. With this context information, the virtual assistant determines that the second candidate intent of obtaining stock price is an actionable intent, because the S&P 500 index information can be obtained.

[0284] FIG. 11B illustrates an exemplary user interface for providing natural language interaction by a virtual assistant (e.g., virtual assistant 1000 shown in FIG. 10) using context

information associated with sensory data. As illustrated in FIGS. 10 and 11B, for example, the virtual assistant operating on device 1120 receives one or more audio streams from user 1110 and user 1130. An audio stream may include, for example, an utterance 1132 from user 1130, an utterance 1134 from user 1130, and an utterance 1136 from user 1110. One or more utterances from user 1110 may be directed to the virtual assistant or user 1130. For example, utterance 1132 from user 1110 may include “Siri, what is the S&P index today?” Similar to those described above, the virtual assistant determines that a first candidate text representation of utterance 1132 includes a lexical trigger (e.g., “Siri”) and therefore is not to be disregarded. Based on such a determination, the virtual assistant generates a first candidate intent and determines that the first candidate intent of obtaining S&P 500 index is actionable. Accordingly, the virtual assistant performs the task of obtaining the S&P 500 index and outputs a result 1138 (e.g., an audio and/or visual output indicating “The S&P 500 Index is at 12000 today.”).

[0285] As shown in FIG. 11B, upon hearing the S&P 500 index outputted by virtual assistant 1000, user 1130 may provide an utterance 1134 asking user 1130 “Did you hear the AAPL went up a lot today?” The virtual assistant operating on device 1120 can generate a second candidate text representation corresponding to utterance 1134 and determine whether the second candidate text representation is to be disregarded. For example, the virtual assistant may determine that the second candidate text representation does not include a lexical trigger, and therefore estimate the likelihood utterance 1134 is not directed to virtual assistant. In some embodiments, to estimate the likelihood, the virtual assistant can obtain sensory data from one or more sensors communicatively coupled to device 1120 and estimate, based on the obtained sensory data, the likelihood that utterance 1134 is not directed to the virtual assistant. For example, device 1120 can include an optical sensor 264 (e.g., a camera) that detects user 1110’s eye gaze with respect to device 1120 at any given time. Optical sensor 264 may detect, for example, that user 1130 is not looking at device 1120 while utterance 1134 is received by the virtual assistant operating on device 1120. As a result, the virtual assistant estimates, using the sensory data provided by optical sensor 264, that the likelihood utterance 1134 is not directed to the virtual assistant is high (e.g., comparing to a threshold). Accordingly, FTM 840 determines the second candidate text representation is to be disregarded.

[0286] Continuing with the example shown in FIG. 11B, upon hearing utterance 1134, user 1130 may provide utterance 1136 “Really, what is the price for AAPL today?” The virtual assistant can generate a third candidate text representation corresponding to utterance 1136 and determine whether the third candidate text representation is to be disregarded. Similarly to those described above, the virtual assistant can determine that the third candidate text representation does not include a lexical trigger. Based on such a determination, the virtual assistant can estimate the likelihood utterance 1136 is not directed to the virtual assistant operating on device 1120. In some embodiments, to estimate the likelihood, the virtual assistant obtains sensory data and estimate, based on the obtained sensory data, the likelihood that utterance 1136 is not directed to the virtual assistant 1000. For example, optical sensor 264 may detect, for example, that user 1110 is looking at device 1120 while utterance 1136 is received by the virtual assistant. As a

result, the virtual assistant estimates, using the sensory data provided by optical sensor **264**, that the likelihood utterance **1136** is not directed to virtual assistant **1000** is low (e.g., comparing to a likelihood threshold). Accordingly, the virtual assistant determines the third candidate text representation is not to be disregarded. The virtual assistant can therefore generate a candidate intent based on the third candidate text representation and determine that candidate intent is an actionable intent. Accordingly, the virtual assistant performs the task to obtain the stock price of AAPL and outputs a result **1140** (e.g., an audio and/or visual output indicating “AAPL closed at \$200 today.”).

[0287] FIGS. 12A-12D illustrates exemplary user interfaces for providing natural language interaction by a virtual assistant operating on device **1220** using context information associated with executing a previously determined actionable intent. As shown in FIGS. 10 and 12A, user **1210** provides a first audio stream **1212** (e.g., “Siri, play some music”), which is received by the virtual assistant operating on device **1220**. Similar to those described above, the virtual assistant determines that first audio stream **1212** includes a lexical trigger (e.g., “Siri”), generates a candidate text representation of first audio stream **1212** (e.g., “Siri, play some music.”), determines that the candidate text representation of first audio stream **1212** is not to be disregarded because it includes the lexical trigger, generates a candidate intent, determines that the candidate intent is an actionable intent (e.g., a task of playing music can be performed), executes the actionable intent, and outputting a result of the execution (e.g., outputting an audio and/or visual message **1224** such as “Here is some music you might like.”). As illustrated in FIG. 12A, virtual assistant **1000** may receive and process first audio stream **1212** at a time **1222A** (e.g., 9 AM).

[0288] As illustrated in FIG. 12B, while the virtual assistant operating on device **1220** is executing the task of playing music, it may receive a second audio stream **1214** from user **1210** at a time **1222B** (e.g., 9:05 AM). In some embodiments, the virtual assistant generates a candidate text representation of the second audio stream **1214**. Based on the candidate text representation of second audio stream **1214**, the virtual assistant can determine whether second audio stream **1214** (e.g., “Stop”) is a part of a same audio session as first audio stream **1212** (e.g., “Play some music”).

[0289] In some embodiments, to determining whether second audio stream **1214** is a part of the same audio session that includes first audio stream **1212**, the virtual assistant operating on device **1220** can obtain context information associated with executing a previously determined actionable intent. For example, the virtual assistant detects that when second audio stream **1214** is received, music is playing as a result of executing the previous determined actionable intent. As a result, the virtual assistant determines that second audio stream **1214** is or is likely a part of the same audio session that includes first audio stream **1212**. In some embodiments, such a determination may also taking into account of a relation (e.g., semantic relation, topical relation) between first audio stream **1212** and second audio stream **1214**. Embodiments that taking into account of a relation is further described in detail below.

[0290] With reference to FIG. 12B, in accordance with a determination that the second audio stream **1214** is a part of the same audio session that includes first audio stream **1212**, the virtual assistant generates, based on the candidate text

representation of second audio stream **1214**, a second candidate intent. The virtual assistant further determines whether the second candidate intent includes an actionable intent. In some embodiments, this determination can be based on context information associated with a previous task performed or is being performed by the virtual assistant. For example, the context information can be associated with performing the previous task of playing music. As shown in FIG. 12B, the second candidate intent may thus be determined to be stopping the music that is currently playing. The virtual assistant can thus determine that a task can be performed (e.g., stop playing the music) according to this second candidate intent, and therefore the second candidate intent is actionable. As a result, the virtual assistant executes the actionable intent to stop playing the music. In some embodiments, the virtual assistant optionally provides an audio and/or visual output **1226** (e.g., “OK”) indicating the results of execution.

[0291] FIGS. 12C and 12D illustrate a similar scenario where the virtual assistant operating on device **1220** receives a first audio stream **1242** including an utterance such as “Siri, play some music.” Similar to those described above, the virtual assistant determines an actionable intent of playing music, executes the actionable intent, and optionally output an audio/visual message **1222** (e.g., “Here is some music you might like.”). While the music is still playing, the virtual assistant receives a second audio stream **1244** including an utterance such as “Skip this song.” The virtual assistant detects that when second audio stream **1244** is received, music is playing as a result of executing the previous determined actionable intent. As a result, the virtual assistant can determine that second audio stream **1244** is or is likely a part of the same audio session that includes first audio stream **1242**. As a result, the virtual assistant generates, based on the candidate text representation of second audio stream **1244**, a second candidate intent. The virtual assistant further determines whether the second candidate intent includes an actionable intent. For example, as shown in FIG. 12D, the second candidate intent may be determined to be skipping sound track **1** that is currently playing and starting to play the next sound track. The virtual assistant can thus determine that a task can be performed according to this second candidate intent, and therefore the second candidate intent is actionable. As a result, the virtual assistant executes the actionable intent to stop playing sound track **1** and start playing sound track **2**.

[0292] FIGS. 13A-13B illustrates exemplary user interfaces for providing natural language interaction by a virtual assistant using context information associated with a relation of user utterances or audio streams. As described above, context information can be provided to one or more modules or components of a virtual assistant (e.g., FTM **840** and CIE **860** of virtual assistant **1000** shown in FIG. 10). In some embodiments, context information can be provided to a CIE (e.g., CIE **860**) for determining whether a candidate intent is an actionable intent. Such context information can include, for example, one or more relations among candidate text representations representing the user utterances in an audio stream.

[0293] As an example shown in FIG. 13A, a virtual assistant (e.g., virtual assistant **1000**) operating on device **1320** receives an audio stream **1312**. Audio stream **1312** may include a first utterance (e.g., “You know warriors did great in their last game.”) and a second utterance (e.g., “When is



the next game, Siri?"). Based on audio stream **1312**, candidate text representations can be generated to represent the utterances included in audio stream **1312**. The virtual assistant can determine (e.g., using FTM **840**) that the candidate text representation of the first utterance (e.g., "You know warriors did great in their last game.") does not include a lexical trigger and that the first utterance is likely not directed to the virtual assistant. As a result, the virtual assistant can determine that the candidate text representation of the first utterance is to be disregarded for the purpose of generating candidate intents. In some embodiments, while the candidate text representation of the first utterance is disregarded for the purpose of generating candidate intents, it can be retained for determining relations among multiple candidate text representations.

[0294] Continuing with the example shown in FIG. **13A**, the virtual assistant can determine (e.g., using FTM **840**) that the candidate text representation of the second utterance (e.g., "When is the next game, Siri?") includes a lexical trigger and thus the second utterance is likely directed to the virtual assistant. As a result, the candidate text representation of the second utterance is not to be disregarded. Accordingly, the virtual assistant generates a candidate intent based on the candidate text representation of the second utterance, and determines whether the candidate intent is an actionable intent. As described above, to determine whether a candidate intent is an actionable intent, in some embodiments, the virtual assistant determines whether a task can be performed according to the actionable intent. The determination of whether a task can be performed can be based on the candidate text representation of the second utterance and based on context information indicating relations among candidate text representations.

[0295] As described above, the candidate text representation of the second utterance in audio stream **1312** may include "When is the next game, Siri?" Based solely on this candidate text representation, the virtual assistant may not be able to determine whether a task can be performed because the candidate text representation does not indicate which game the user refers to. In some examples, the virtual assistant can determine relations among the adjacent candidate text representations. For example, the virtual assistant determines that the candidate text representation of the first utterance (e.g., "You know warriors did great in their last game.") is semantically, topically, and/or temporally related to the candidate text representation of the second utterance (e.g., "When is the next game, Siri?"). As a result, the virtual assistant can use context information associated with the relations among the candidate text representations to determine whether a task can be performed. In this example, the context information indicates that user **1310** is likely referring to the Golden State Warriors game when user **1310** asks "When is the next game, Siri?" As a result, the virtual assistant determines that a task can be performed based on the context information (e.g., a task to obtain the schedule of the Warriors game can be performed by searching the Internet) and thus determines the candidate intent is an actionable intent. Accordingly, the virtual assistant can execute the actionable intent and provides, for example, an audio and/or visual output **1326** such as "Golden State Warriors next game is schedule tomorrow at 5:30 pm."

[0296] Turning now to FIG. **13B**, context information associated with relations (e.g., semantic, topical, temporal) among multiple audio streams can also be used to determine

whether two or more audio streams are in the same audio session. With reference to FIG. **13B**, the virtual assistant (e.g., virtual assistant **1000** shown in FIG. **10**) operating on device **1320** receives a first audio stream **1342** including an utterance such as "What was the score of the Warriors' game?" Similar to those described above, the virtual assistant generates a candidate text representation of first audio stream **1342**, determines that it is not to be disregarded, generates a candidate intent, determines that the candidate intent is an actionable intent, executes the actionable intent, and provides an audio and/or visual output **1346** such as "Warriors defeated Knicks by a score of 123 to 112."

[0297] In some embodiments, while or after the virtual assistant outputs a result of execution of the actionable intent corresponding to first audio stream **1342**, the virtual assistant operating on device **1320** receives a second audio stream **1344** including an utterance such as "When is the next game?" In some embodiments, the virtual assistant generates a candidate text representation to represent second audio stream **1344**. The generation of the candidate text representation of second audio stream **1344** can be performed irrespective of whether second audio stream **1344** includes a lexical trigger. For example, the virtual assistant can determine a temporal relation between receiving second audio stream **1344** and the execution of the actionable intent corresponding to first audio stream **1342**. The virtual assistant further determines that the temporal relation indicates that second audio stream **1344** is received within a threshold time period from the execution of the actionable intent, and therefore a lexical trigger in second audio stream **1344** is not required for generating a candidate text representation of second audio stream **1344**.

[0298] After generation of the candidate text representation of second audio stream **1344**, the virtual assistant operating on device **1320** can determine whether second audio stream **1344** is a part of a same audio session that includes first audio stream **1342**. The virtual assistant can determine relations among the respective candidate text representations of first audio stream **1322** and second audio stream **1344**. For example, the virtual assistant determines that the candidate text representation of the first audio stream **1322** (e.g., "What was the score of the Warriors game?") is semantically, topically, and/or temporally related to the candidate text representation of the second audio stream (e.g., "When is the next game?"). As a result, the virtual assistant can use context information associated with the relations among the respective candidate text representations of the first and second audio streams to determine whether the audio streams are in the same audio session. In the example shown in FIG. **13B**, the virtual assistant determines that second audio stream **1344** is a part of the same audio session that include first audio stream **1322**, because the respective candidate text representations of the audio streams are semantically and/or topically related. In accordance with such a determination, the virtual assistant generates a candidate intent based on the candidate text representations of second audio stream **1322**, determines whether the candidate intent is an actionable intent, and if so, executes the actionable intent. As a result, the virtual assistant provides an audio and/or visual output **1348** such as "Golden State Warriors next game is scheduled tomorrow, at 5:30 pm."

[0299] FIGS. **14A-14D** illustrate exemplary user interfaces for selecting a task from a plurality of tasks using

context information. As described above, previous execution results and/or context information can be provided to one or more modules or components of a virtual assistant (e.g., FTM 840 and CIE 860 of virtual assistant 1000 shown in FIG. 10). In some embodiments, context information can be provided to task execution module 880 for the purpose of selecting one or more tasks to perform. With reference to FIG. 14A, a virtual assistant operating on device 1420 receives a first audio stream 1412 including an utterance such as “Siri, play some music.” Similar to those described above with respect to FIG. 12A, the virtual assistant determines that audio stream 1412 includes a lexical trigger, generates a candidate text representation, determines that the candidate text representation is not to be disregarded, generates a candidate intent, determines that the candidate intent is an actionable intent, and executes the actionable intent. As a result, the virtual assistant starts to play music and optionally provides an audio and/or visual output 1424 such as “Here is some music you might like.”

[0300] With reference to FIG. 14B, in some embodiments, while the virtual assistant is executing the previously determined actionable intent (e.g., playing music), it receives a second audio stream 1414. Second audio stream 1414 may include another user request, which may not semantically or topically relate to first audio stream 1412. For example, second audio stream 1414 includes an utterance such as “I would like to send a message to my wife, Siri.” The virtual assistant operating on device 1420 can repeat a process similar to those described above and execute a second actionable intent derived based on second audio stream 1414. As a result, the virtual assistant causes device 1420 to display, for example, a text message user interface 1426 for composing a text message.

[0301] Turning to FIG. 14C, in some embodiments, while executing both of the previously determined actionable intents (e.g., playing music and displaying a text message user interface for composing a text message), the virtual assistant operating on device 1420 receives a third audio stream 1416. Third audio stream 1416 may be related to the execution of one or both of the previously determined actionable intents. For example, third audio stream 1416 may include an utterance such as “stop.” Based on third audio stream 1416, the virtual assistant can determine, for example, a first candidate intent of stopping the currently-playing music and a second candidate intent of stopping composing of the text message. The virtual assistant can further determine that both the first and second candidate intents are actionable intents because a respective task can be performed.

[0302] In some embodiments, if there is a plurality of actionable intents, the virtual assistant selects, from a plurality of tasks associated with the plurality of actionable intents, a single task to perform. In some embodiments, the selection of a single task for execution can be based on context information. For example, the virtual assistant obtains context information associated with a most-recent task initiated by the virtual assistant and selects a single task for execution based on such context information. In this example as described above, the most-recent task initiated by the virtual assistant is to display a user interface for composing a text message. Accordingly, based on the context information associated with this most-recent task of displaying user interface for composing a text message, the virtual assistant disambiguates the user request included in

third audio stream 1416 and selects the task of stopping or cancelling the composing of the text message. Optionally, as shown in FIG. 14C, the virtual assistant can provide an audio and/or visual output 1432 such as “OK, message to your wife cancelled.”

[0303] In some embodiments, if there is a plurality of actionable intents, the virtual assistant selects, from a plurality of tasks associated with the plurality of actionable intents, a single task to perform based on a user selection. This is illustrated in FIG. 14D, where the virtual assistant outputs a plurality of task options. For example, the virtual assistant provides an audio and/or visual output 1434 (e.g., “Would you like to stop the music or cancel the message to your wife?”) prompting user 1410 to select a task. The virtual assistant receives a user selection in an audio stream 1418 (e.g., “stop the music”) and performs the task based on the user selection. Optionally, as shown in FIG. 14D, the virtual assistant can provide an audio and/or visual output 1436 such as “Music stopped.”

[0304] In some embodiments, if there is a plurality of actionable intents, the virtual assistant selects, from a plurality of tasks associated with the plurality of actionable intents, a single task to perform based on a priority associated with each of the plurality of tasks. For example, a virtual assistant (e.g., virtual assistant 1000) may receive a first audio stream including an utterance such as “How is the weather today, Siri?” Before the virtual assistant responds, it receives a second audio stream including utterances such as “There is a car accident! Siri, call 911!” Based on the audio streams, the virtual assistant determines a first actionable intent of obtaining weather information and a second actionable intent of making an emergency call. In some embodiments, the virtual assistant can determine a priority associated with each of the plurality of tasks to be performed and select the task based on the determined priorities. In the above example, while the second audio stream is received after receiving the first audio stream, the virtual assistant determines that the task associated with the second actionable intent of making an emergency 911 call has a higher priority than the task associated with the first actionable intent of obtaining weather information. Accordingly, the virtual assistant selects the task associated with the second actionable intent for execution (e.g., making an emergency 911 call).

[0305] While the various embodiments described above relate to specific type of context information, it is appreciated that the techniques described can also use any types of context information, as described in U.S. patent application Ser. No. 15/694,267, “Methods and Systems for Customizing Suggestions Using User-specific Information,” filed Sep. 1, 2017, which is hereby incorporated by reference in its entirety.

## 5. Process for Providing Natural Language Interaction

[0306] FIG. 15A-15G illustrates process 1500 for operating a virtual assistant for providing natural language interaction, according to various embodiments. Process 1500 is performed, for example, using one or more electronic devices implementing a virtual assistant. In some examples, process 1500 is performed using a client-server system (e.g., system 100), and the blocks of process 1500 are divided up in any manner between the server (e.g., DA server 106) and a client device. In other examples, the blocks of process 1500 are divided up between the server and multiple client

devices (e.g., a mobile phone and a smart watch). Thus, while portions of process **1500** are described herein as being performed by particular devices of a client-server system, it will be appreciated that process **1500** is not so limited. In other examples, process **1500** is performed using only a client device (e.g., user device **104**, **200**, **400**, **600**, **900**, **1120**, **1220**, **1320**, or **1420**) or only multiple client devices. In process **1500**, some blocks are, optionally, combined, the order of some blocks is, optionally, changed, and some blocks are, optionally, omitted. In some examples, additional steps may be performed in combination with the process **1500**.

[**0307**] As described above, always requiring a trigger phrase at the beginning portion of an utterance from the user can cause the human-machine interaction to become cumbersome and make the human-machine user interface less natural and efficient. The techniques described in this application, including those represented by process **1500**, eliminate or reduce the need of this requirement to lead every user utterance with a trigger phrase. Instead, a trigger word or phrase can be placed in any portion of an audio stream that may include one or more user utterances. Moreover, the techniques described in this application do not require using a trigger phrase that include a plurality of words (e.g., “Hey Siri”). A single word (e.g., “Siri”) can be used to indicate that the audio stream including the user utterances is directed to the virtual assistant. This enables a more natural way of communication. As a result, the techniques enhance the operability of the device and make the user-device interface more efficient which, additionally, reduces power usage and improves battery life of the device by enabling the user to use the device more quickly and efficiently.

[**0308**] With reference to FIG. **15A**, at block **1502**, first audio stream (e.g., audio stream **912** as illustrated in FIG. **9**) including one or more utterances is received via a microphone. At block **1504**, whether the first audio stream includes a lexical trigger is determined. In some embodiments, the lexical trigger is a single-word lexical trigger. In some embodiments, the first audio stream includes a first utterance and the single-word lexical trigger is positioned in a portion of the first utterance other than the beginning portion of the first utterance. As an example described in FIG. **9**, the lexical trigger (e.g., “Siri”) is positioned at the end of an utterance (e.g., “Turn on the light, Siri.”).

[**0309**] At block **1506**, to determine whether the first audio stream includes a lexical trigger, a beginning point of the first audio stream is detected. As one example of detecting the beginning point of the first audio stream, at block **1508**, an absence of voice activity before receiving the first audio stream is detected. At block **1510**, whether the absence of voice activity before receiving the first audio stream exceeds a first threshold period of time is determined. At block **1512**, in accordance with a determination that the absence of voice activity exceeds the first threshold period of time, the beginning point of the first audio stream is determined based on the absence of voice activity before receiving the first audio stream.

[**0310**] At block **1514**, to determine whether the first audio stream includes a lexical trigger, an end point of the first audio stream is detected. As one example of detecting the end point of the first audio stream, at block **1516**, an absence of voice activity after receiving the one or more utterances of the first audio stream is detected via the microphone. At block **1518**, whether the absence of voice activity after

receiving the one or more utterances of the first audio stream exceeds a second threshold period of time is determined. At block **1520**, in accordance with a determination that the absence of voice activity after receiving the one or more utterances of the first audio stream exceeds the second threshold period of time, the end point of the first audio stream is determined based on the absence of voice activity after receiving the one or more utterances of the first audio stream.

[**0311**] As another example of detecting the end point of the first audio stream, at block **1522**, a pre-configured duration that the electronic device is configured to receive the first audio stream is obtained. With reference to FIG. **15B**, at block **1524**, the end point of the first audio stream is determined based on the detected beginning point of the first audio stream and the pre-configured duration.

[**0312**] As another example of detecting the end point of the first audio stream, at block **1526**, a size of an audio file representing the received one or more utterances of the first audio stream is determined. At block **1528**, the size of the audio file is compared with a capacity of a buffer storing the audio file. At block **1530**, the end point of the first audio stream is determined based on a result of comparing the size of the audio file with the capacity of the buffer storing the audio file. Various embodiments of detecting of the beginning point and the end point of an audio stream are described in detail above with respect to, for example, FIGS. **8** and **9**.

[**0313**] At block **1532**, to determine whether the first audio stream includes a lexical trigger (e.g., “Siri”), whether a lexical trigger is included between the beginning point and the end point of the first audio stream is determined.

[**0314**] At block **1534**, in accordance with a determination that the first audio stream includes the lexical trigger, one or more candidate text representations of the one or more utterances are generated. As described above with respect to FIG. **8**, candidate text representations can be generated by a natural language engine (e.g., NLE **820**).

[**0315**] At block **1536**, to generate the one or more candidate text representations, speech-to-text conversion of each of the one or more utterances of the first audio stream to generate the one or more candidate text representations is performed. For example, speech-to-text conversion of each of the utterances of audio stream **912** shown in FIG. **9** can be performed to generate candidate text representations including a first candidate text representation (e.g., “It’s dark outside.”) and a second candidate text representation (e.g., “Turn on the light, Siri.”). At block **1538**, confidence levels corresponding to the one or more candidate text representations are determined.

[**0316**] With reference to FIG. **15C**, at block **1540**, whether at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant is determined. At block **1542**, to determine whether at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant, whether the at least one candidate text representation includes the lexical trigger is determined. At block **1544**, in accordance with a determination that the at least one candidate text representation does not include the lexical trigger, a likelihood that the utterance corresponding to the at least one candidate text representation is not directed to the virtual assistant is estimated. As described above, in the example shown in FIG. **9**, an FTM of the virtual assistant operating on device **900** determines that a

candidate text representation (e.g., “It’s dark outside.”) does not include a lexical trigger. Further, using a decision tree, the FTM of the virtual assistant can estimate a likelihood that the utterance corresponding to the particular first candidate text representation is not directed to the virtual assistant.

[0317] As one example of estimating the likelihood that the utterance corresponding to the at least one candidate text representation is not directed to the virtual assistant, at block 1546, context information associated with a usage pattern (e.g., frequency of a particular question is asked at a particular time, as illustrated in FIGS. 10 and 11A) of the virtual assistant is obtained. At block 1548, based on the context information associated with the usage pattern of the virtual assistant, the likelihood that the utterance corresponding to the at least one candidate text representation is not directed to the virtual assistant is estimated.

[0318] As another example of estimating the likelihood that the utterance corresponding to the at least one candidate text representation is not directed to the virtual assistant, at block 1550, sensory data (e.g., location data as described above) from one or more sensors communicatively coupled to the electronic device are obtained. At block 1552, based on sensory data, the likelihood that the utterance corresponding to the at least one candidate text representation is not directed to the virtual assistant is estimated.

[0319] As another example of estimating the likelihood that the utterance corresponding to the at least one candidate text representation is not directed to the virtual assistant, a degree that a candidate text representation conforms to a language model (LM) and/or a context free grammar (CFG) corresponding to recognized/valid requests for virtual assistant is determined. For example, a candidate text representation such as “Eat your vegetables” may have a low degree of conformation to an LM and/or CFG, and is thus not less likely directed to the virtual assistant. Another candidate text representation such as “Book a table for two” may have a high degree of conformation to an LM and/or CFG, and is thus more likely directed to the virtual assistant.

[0320] At block 1554, based on the estimated likelihood, whether the at least one candidate text representation of the one or more candidate text representations is to be disregarded by the virtual assistant is determined. As one example described above with respect to FIG. 11A, based on the context information indicating that the user frequently asks the virtual assistant about the stock price around 6 AM in the morning and that the current time is about 6 AM, the virtual assistant can estimate that the likelihood the utterance “And do you know what is the stock price?” is not directed to the virtual assistant is low (e.g., does not satisfy a threshold). As a result, the virtual assistant determines, based on the estimated likelihood, that the particular candidate text representation is not to be disregarded.

[0321] With reference to FIG. 15D, at block 1556, in accordance with a determination that at least one candidate text representation is to be disregarded by the virtual assistant, one or more candidate intents are generated based on candidate text representations of the one or more candidate text representations other than the to be disregarded at least one candidate text representation. To generate the candidate intents, at block 1558, one or more pre-mitigation intents corresponding to the one or more candidate text representations of the one or more utterances are obtained. At block 1560, from the one or more pre-mitigation intents, the one

or more candidate intents corresponding to the one or more candidate text representations other than the to be disregarded at least one candidate text representation are selected. As described above with respect to FIG. 10, candidate intents can be generated by an FTM (e.g., FTM 1044).

[0322] At block 1562, whether the one or more candidate intents include at least one actionable intent is determined. At block 1564, to determine whether the one or more candidate intents include at least one actionable intent, for each of the one or more candidate intents, whether a task can be performed is determined. As described above, in some embodiments, such a determination can be performed by a CIE (e.g., CIE 860 shown in FIG. 10). As one example of determining whether a task can be performed, at block 1566, context information associated with a usage pattern of the virtual assistant is obtained. At block 1568, based on the context information associated with the usage pattern of the virtual assistant, whether the task can be performed is determined. One example of such a determination based on usage pattern is illustrated above with respect to FIG. 11B.

[0323] As another example of determining whether a task can be performed, at block 1570, context information associated with a previous task performed by the virtual assistant is obtained. At block 1572, based on the context information associated with the previous task performed by the virtual assistant, whether the task can be performed is determined. One example of such a determination based on previous task performed by the virtual assistant is illustrated above with respect to FIGS. 12A-12D.

[0324] With reference to FIG. 15E, as another example of determining whether a task can be performed, at block 1574, one or more relations among the one or more candidate text representations is determined. At block 1576, whether the task can be performed based on the one or more relations among the one or more candidate text representations is determined. One example of such a determination based on one or more relations is illustrated above with respect to FIGS. 13A-13B.

[0325] As another example of determining whether a task can be performed, at block 1578, sensory data (e.g., location data) from one or more sensors communicatively coupled to the electronic device are obtained. At block 1580, whether the task can be performed based on the sensory data is determined.

[0326] As another example of determining whether the task can be performed, at block 1582, a confidence level associated with performing the task is estimated. At block 1584, whether the confidence level associated with performing the task satisfies a threshold confidence level is determined. At block 1586, in accordance with a determination that the confidence level associated with performing the task satisfies the threshold confidence level, it is determined that the task can be performed.

[0327] At block 1588, in accordance with a determination that the task can be performed, it is determined that the one or more candidate intents include at least one actionable intent. For example, as described above with respect to FIGS. 11A, based on context information (e.g., a usage pattern and a time), the virtual assistant determines that a particular candidate intent of obtaining stock price is an actionable intent.

[0328] At block 1590, in accordance with a determination that the one or more candidate intents include at least one actionable intent, the at least one actionable intent is

executed. At block **1592**, one or more tasks are performed according to the at least one actionable intent.

[0329] With reference to FIG. **15F**, to execute the at least one actionable intent, at block **1594**, a first task for execution is selected from a plurality of tasks associated with the plurality of actionable intents. As one example of selecting the first task for execution, at block **1596**, context information associated with a most-recent task initiated by the virtual assistant (e.g., displaying a user interface for composing a text message as illustrated in FIG. **14B**) is obtained. At block **1598**, the first task is selected based on the context information associated with a previous task performed by the virtual assistant.

[0330] As one example of selecting the first task for execution, at block **1600**, a plurality of task options corresponding to the plurality of tasks associated with the plurality of actionable intents is outputted. At block **1602**, a user selection is received from the plurality of task options. At block **1604**, the first task for execution is selected based on the user selection. Selecting a task for execution based on a user selection is illustrated above with respect to FIG. **14D**.

[0331] As one example of selecting the first task for execution, at block **1606**, a priority associated with each of the plurality of tasks is determined. At block **1608**, the first task for execution is selected based on the priority associated with each of the plurality of tasks. At block **1610**, the selected first task is performed. As described above, a task having a higher priority (e.g., making an emergency call) is selected over a task having a lower priority (e.g., reporting weather information).

[0332] At block **1612**, a result of the execution of the at least one actionable intent is outputted.

[0333] With reference to FIG. **15G**, at block **1614**, upon executing the at least one actionable intent, a second audio stream is received via the microphone. At block **1616**, one or more second candidate text representations are generated to represent the second audio stream. At block **1618**, based on the one or more second candidate text representations, whether the second audio stream is a part of an audio session that includes the first audio stream is determined.

[0334] At block **1620**, as one example of determining whether the second audio stream is a part of an audio session that includes the first audio stream, context information associated with executing the at least one actionable intent is obtained. At block **1622**, based on the context information associated with executing the at least one actionable intent, whether the second audio stream is a part of the audio session that includes the first audio stream is determined.

[0335] At block **1624**, as another example of determining whether the second audio stream is a part of an audio session that includes the first audio stream, a relation among respective candidate text representations of the first audio stream and the second audio stream is determined. At block **1626**, based on the relation among respective candidate text representations of the first audio stream and the second audio stream, whether the second audio stream is a part of the audio session that includes the first audio stream is determined.

[0336] At block **1628**, in accordance with a determination that the second audio stream is a part of the audio session that includes the first audio stream, one or more second candidate intents are generated based on the one or more second candidate text representations.

[0337] At block **1630**, whether the one or more second candidate intents include at least one second actionable intent is determined. At block **1632**, in accordance with a determination that the one or more second candidate intents include at least one second actionable intent, the at least one second actionable intent is executed. At block **1634**, a result of the execution of the at least one second actionable intent is outputted. Examples of the process illustrated in FIG. **15G** are described above with respect to FIGS. **12A-12D**.

[0338] The foregoing description, for purpose of explanation, has been described with reference to specific embodiments. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The embodiments were chosen and described in order to best explain the principles of the techniques and their practical applications. Others skilled in the art are thereby enabled to best utilize the techniques and various embodiments with various modifications as are suited to the particular use contemplated.

[0339] Although the disclosure and examples have been fully described with reference to the accompanying drawings, it is to be noted that various changes and modifications will become apparent to those skilled in the art. Such changes and modifications are to be understood as being included within the scope of the disclosure and examples as defined by the claims.

[0340] As described above, one aspect of the present technology is the gathering and use of data available from various sources (e.g., the context information associated with a usage pattern of the virtual assistant or device) to improve human-machine interface to provide a more natural language interaction. The present disclosure contemplates that in some instances, this gathered data may include personal information data that uniquely identifies or can be used to contact or locate a specific person. Such personal information data can include demographic data, location-based data, telephone numbers, email addresses, twitter IDs, home addresses, data or records relating to a user's health or level of fitness (e.g., vital signs measurements, medication information, exercise information), date of birth, or any other identifying information.

[0341] The present disclosure recognizes that the use of such personal information data, in the present technology, can be used to the benefit of users. For example, the personal information data can be used to deliver targeted content that is of greater interest to the user. Accordingly, use of such personal information data enables calculated control of the delivered content. Further, other uses for personal information data that benefit the user are also contemplated by the present disclosure. For instance, health and fitness data may be used to provide insights into a user's general wellness, or may be used as positive feedback to individuals using technology to pursue wellness goals.

[0342] The present disclosure further contemplates that the entities responsible for the collection, analysis, disclosure, transfer, storage, or other use of such personal information data will comply with well-established privacy policies and/or privacy practices. In particular, such entities should implement and consistently use privacy policies and practices that are generally recognized as meeting or exceeding industry or governmental requirements for maintaining personal information data private and secure. Such policies should be easily accessible by users, and should be updated

as the collection and/or use of data changes. Personal information from users should be collected for legitimate and reasonable uses of the entity and not shared or sold outside of those legitimate uses. Further, such collection/sharing should occur only after receiving the informed consent of the users. Additionally, such entities would take any needed steps for safeguarding and securing access to such personal information data and ensuring that others with access to the personal information data adhere to their privacy policies and procedures. Further, such entities can subject themselves to evaluation by third parties to certify their adherence to widely accepted privacy policies and practices. In addition, policies and practices should be adapted for the particular types of personal information data being collected and/or accessed and adapted to applicable laws and standards, including jurisdiction-specific considerations. For instance, in the US, collection of or access to certain health data may be governed by federal and/or state laws, such as the Health Insurance Portability and Accountability Act (HIPAA); whereas health data in other countries may be subject to other regulations and policies and should be handled accordingly. Hence different privacy practices should be maintained for different personal data types in each country.

**[0343]** Despite the foregoing, the present disclosure also contemplates embodiments in which users selectively block the use of, or access to, personal information data. That is, the present disclosure contemplates that hardware and/or software elements can be provided to prevent or block access to such personal information data. For example, in the case of collecting usage pattern of a user's activities, the present technology can be configured to allow users to select to "opt in" or "opt out" of participation in the collection of personal information data before or during such a collection. In another example, users can select not to provide or share the users' activities information. In yet another example, users can select to limit the length of time the users' activities information is maintained or entirely prohibit the development of the usage pattern based on the activities information. In addition to providing "opt in" and "opt out" options, the present disclosure contemplates providing notifications relating to the access or use of personal information. For instance, a user may be notified upon downloading an app that their personal information data will be accessed and then reminded again just before personal information data is accessed by the app.

**[0344]** Moreover, it is the intent of the present disclosure that personal information data should be managed and handled in a way to minimize risks of unintentional or unauthorized access or use. Risk can be minimized by limiting the collection of data and deleting data once it is no longer needed. In addition, and when applicable, including in certain health related applications, data de-identification can be used to protect a user's privacy. De-identification may be facilitated, when appropriate, by removing specific identifiers (e.g., date of birth, etc.), controlling the amount or specificity of data stored (e.g., collecting location data at a city level rather than at an address level), controlling how data is stored (e.g., aggregating data across users), and/or other methods.

**[0345]** Therefore, although the present disclosure broadly covers use of personal information data to implement one or more various disclosed embodiments, the present disclosure also contemplates that the various embodiments can also be

implemented without the need for accessing such personal information data. That is, the various embodiments of the present technology are not rendered inoperable due to the lack of all or a portion of such personal information data. For example, the likelihood that a user utterance is directed to a virtual assistant can be estimated based on non-personal information data or a bare minimum amount of personal information, such as the content being requested by the device associated with a user, other non-personal information available to the virtual assistant, or publically available information.

What is claimed is:

1. An electronic device, comprising:
  - one or more processors;
  - a microphone; and
  - memory storing one or more programs configured to be executed by the one or more processors, the one or more programs including instructions for:
    - providing a spoken output;
    - receiving a user utterance;
    - determining whether the user utterance references the spoken output;
    - in response to determining that the user utterance references the spoken output:
      - obtaining user gaze data from a sensor communicatively coupled to the electronic device;
      - determining, based on the user gaze data, a likelihood that the user utterance is directed to a virtual assistant;
      - determining whether the likelihood that the user utterance is directed to the virtual assistant exceeds a threshold; and
      - in accordance with a determination that the likelihood that the user utterance is directed to the virtual assistant exceeds a threshold, determining a candidate intent based on the user utterance.
2. The electronic device of claim 1, wherein the sensor is a first sensor and wherein the one or more programs further include instructions for:
  - obtaining context information from a second sensor communicatively coupled to the electronic device, wherein the likelihood that the user utterance is directed to the virtual assistant is further based on the context information.
3. The electronic device of claim 1, wherein the one or more programs further include instructions for:
  - in accordance with a determination that the likelihood that the user utterance is directed to the virtual assistant is below a threshold, disregarding the user utterance.
4. The electronic device of claim 1, wherein the one or more programs further include instructions for:
  - determining whether a task associated with the candidate intent can be performed.
5. The electronic device of claim 4, wherein the one or more programs further include instructions for:
  - in accordance with a determination that the task associated with the candidate intent can be performed:
    - performing the task; and
    - providing an output indicative of the task.
6. The electronic device of claim 1, wherein likelihood that the user utterance is directed to a virtual assistant is high when the user gaze is directed towards the electronic device.

7. The electronic device of claim 1, wherein the user utterance is a first user utterance and wherein the one or more programs further include instructions for:

in accordance with a determination that the likelihood that the user utterance is directed to the virtual assistant is below a threshold:

receiving a second user utterance;

obtaining user gaze data from the sensor communicatively coupled to the electronic device;

determining, based on the user gaze data, a likelihood that the second user utterance is directed to a virtual assistant;

in accordance with a determination that the likelihood that the second user utterance is directed to the virtual assistant exceeds a threshold, determining a candidate intent based on the user utterance.

8. The electronic device of claim 1, wherein the one or more programs further include instructions for:

determining whether the user utterance includes a lexical trigger, wherein the user gaze data is obtained in response to a determination that the user utterance does not include a lexical trigger.

9. A method for providing natural language interaction by a virtual assistant, the method comprising:

at an electronic device with one or more processors, memory, and a microphone:

providing a spoken output;

receiving a user utterance;

determining whether the user utterance references the spoken output;

in response to determining that the user utterance references the spoken output:

obtaining user gaze data from a sensor communicatively coupled to the electronic device;

determining, based on the user gaze data, a likelihood that the user utterance is directed to a virtual assistant;

determining whether the likelihood that the user utterance is directed to the virtual assistant exceeds a threshold; and

in accordance with a determination that the likelihood that the user utterance is directed to the virtual assistant exceeds a threshold, determining a candidate intent based on the user utterance.

10. A non-transitory computer-readable storage medium storing one or more programs configured to be executed by one or more processors of an electronic device, the one or more programs including instructions for:

provide a spoken output;

receive a user utterance;

determine whether the user utterance references the spoken output;

in response to determining that the user utterance references the spoken output:

obtain user gaze data from a sensor communicatively coupled to the electronic device;

determine, based on the user gaze data, a likelihood that the user utterance is directed to a virtual assistant;

determine whether the likelihood that the user utterance is directed to the virtual assistant exceeds a threshold; and

in accordance with a determination that the likelihood that the user utterance is directed to the virtual assistant exceeds a threshold, determine a candidate intent based on the user utterance.

\* \* \* \* \*