



- (51) **International Patent Classification:**
G10L 13/00 (2006.01)
- (21) **International Application Number:**
PCT/US2014/040393
- (22) **International Filing Date:**
30 May 2014 (30.05.2014)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
61/832,753 7 June 2013 (07.06.2013) US
- (71) **Applicant: APPLE INC.** [US/US]; 1 Infinite Loop, Cupertino, CA 95014 (US).
- (72) **Inventors; and**
- (71) **Applicants : NAIK, Devang, K.** [US/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). **WEINER, Liam** [US/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). **BINDER, Justin, G.** [US/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). **SRISUWANANUKORN, Charles** [US/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). **EVERMANN, Gunnar** [DE/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US). **WILLIAMS, Shaun** [US/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA

95014 (US). **CHEN, Hong** [CN/US]; c/o Apple Inc., 1 Infinite Loop, Cupertino, CA 95014 (US). **NAPOLITANO, Lia, T.** [US/US]; c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US).

(72) **Inventor: GRUBER, Thomas, R.;** c/o Apple Inc., 1 Infinite Loop, Mail Stop 39-SIR1, Cupertino, CA 95014 (US).

(74) **Agents: ED3E, Christopher, B.** et al; Morrison & Foerster LLP, 755 Page Mill Road, Palo Alto, CA 94304-1018 (US).

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

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

[Continued on nextpage]

(54) **Title:** SYSTEM AND METHOD FOR USER-SPECIFIED PRONUNCIATION OF WORDS FOR SPEECH SYNTHESIS AND RECOGNITION




FIG. 4

(57) **Abstract:** The method is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. A first speech input including at least one word is received. A first phonetic representation of the at least one word is determined, the first phonetic representation comprising a first set of phonemes selected from a speech recognition phonetic alphabet. The first set of phonemes is mapped to a second set of phonemes to generate a second phonetic representation, where the second set of phonemes is selected from a speech synthesis phonetic alphabet. The second phonetic representation is stored in association with a text string corresponding to the at least one word.



SYSTEM AND METHOD FOR USER-SPECIFIED PRONUNCIATION OF WORDS FOR SPEECH SYNTHESIS AND RECOGNITION

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority from U.S. Provisional Serial No. 61/832,753, filed on June 7, 2013, entitled SYSTEM AND METHOD FOR USER-SPECIFIED PRONUNCIATION OF WORDS FOR SPEECH SYNTHESIS AND RECOGNITION, which is hereby incorporated by reference in its entirety for all purposes.

TECHNICAL FIELD

[0002] The disclosed implementations relate generally to digital assistants, and more specifically, to digital assistants that make use of user-specified pronunciations of words for speech synthesis and recognition.

BACKGROUND

[0003] Just like human personal assistants, digital assistants or virtual assistants can perform requested tasks and provide requested advice, information, or services. An assistant's ability to fulfill a user's request is dependent on the assistant's correct comprehension of the request or instructions. Recent advances in natural language processing have enabled users to interact with digital assistants using natural language, in spoken or textual forms, rather than employing a conventional user interface (e.g., menus or programmed commands). Such digital assistants can interpret the user's input to infer the user's intent, translate the inferred intent into actionable tasks and parameters, execute operations or deploy services to perform the tasks, and produce outputs that are intelligible to the user. Ideally, the outputs produced by a digital assistant should fulfill the user's intent expressed during the natural language interaction between the user and the digital assistant.

[0004] Digital assistants that interact with users via speech inputs and outputs typically employ speech-to-text processing techniques to convert speech inputs to textual forms that can be further processed, and speech synthesis techniques to convert textual outputs to speech. In both cases, accurate conversion between speech and text is important to the usefulness of the digital assistant. For example, if the words in a speech input are

incorrectly identified by a speech-to-text process, the digital assistant may not be able to properly infer the user's intent, or may provide incorrect or unhelpful responses. On the other hand, if the words in a speech output are incorrectly pronounced by the digital assistant, the user may have difficulty understanding the digital assistant. Moreover, incorrect pronunciations by the digital assistant make the assistant seem less polished and less capable, and may reduce users' interest and confidence in the digital assistant.

[0005] For many words, accurate recognition and synthesis are relatively easy, because their pronunciations are fairly standard, at least between people with similar accents or from similar geographical regions. However, certain words or classes of words may be subject to many different pronunciations, making accurate recognition and synthesis more difficult. For example, proper names are often subject to different pronunciations by different people, and it is often not possible to discern the correct pronunciation based only on the spelling of the name. This ambiguity in the correct (or preferred) pronunciation of names is a possible source of recognition and synthesis errors by a digital assistant.

[0006] Accordingly, there is a need for systems and methods to allow users to specify pronunciations of words for recognition and synthesis by a digital assistant.

SUMMARY

[0007] The implementations described herein allow users to specify a correct or a preferred pronunciation of words, such as proper names, so that a digital assistant can both recognize and synthesize the word based on the specified pronunciation. (The term "correct" does not necessarily refer to a single, universally correct pronunciation, as several different pronunciations of a word may be considered "correct" by various individuals and/or in various dialects, accents, languages, etc.) For example, a user can speak a word or name to a digital assistant, and the digital assistant will adjust its speech recognition process to associate the specified pronunciation with the word, and adjust its speech synthesis process so that, when the word is spoken by the digital assistant, it conforms to the user-specified pronunciation. In many cases, however, speech recognition and speech synthesis are performed by different processes using different phonetic representations of the word. Specifically, speech recognition phonetic alphabets are often different from speech synthesis phonetic alphabets, because the phonemes used to generate speech are typically different from those used to recognize speech. More specifically, a speech recognizer may not be able to (or may not need to) detect as many phonemes as a speech synthesizer can produce. And

though a digital assistant can detect phonemes in a speech input in order to learn a user-specified pronunciation, those phonemes may not be suitable for use by a speech synthesizer. Accordingly, separate phonetic representations of the word must be used for speech recognition and synthesis processes.

[0008] Some implementations described herein generate phonetic representations for both speech recognition and synthesis based on a single spoken input. By using only a single spoken input to train speech recognition and speech synthesis processes, the number of interactions necessary to train the digital assistant can be reduced, making the digital assistant appear smarter and more human. Moreover, accepting a spoken input instead of requiring the user to type or otherwise select a textual phonetic representation in a phonetic alphabet allows a more human-like interaction with the digital assistant, thus enhancing the user experience and potentially increasing the user's confidence in the capabilities of the digital assistant.

[0009] Using a single speech input also offers several benefits over techniques that require a user to type in or otherwise select textual phonetic representations of a word. For example, users may be unfamiliar with the particular phonetic alphabet used to train the digital assistant. And if the textual phonetic representations are simplified so that users can use a standard alphabet (e.g., the Latin alphabet) to provide a phonetic representation, differences in accents may result in further confusion and lead to the selection of incorrect pronunciations. As a specific example, for the name "Philippe," a user may represent the first syllable as "fill" in an effort to teach the digital assistant the preferred pronunciation. But even the word "fill" may be pronounced differently by different people (e.g., a person with a French accent may understand "fill" to be pronounced similar to "feel"). Thus, attempting to specify word pronunciations using textual inputs often fail to solve, and can even increase, pronunciation errors. Accordingly, accepting a spoken input of a word simplifies the process of teaching a digital assistant how to recognize and synthesize a word, and increases the accuracy of the teaching process.

[0010] Also described herein are techniques whereby a digital assistant detects when its assumption about a word or name pronunciation is incorrect, and engages in a conversation with the user to acquire the correct pronunciation. For example, if a user asks a digital assistant to "send a text message to Philippe" (pronouncing the name "fill-eep-ay") the digital assistant may not properly recognize the name, and may tell the user that it doesn't understand the input. The user may then bypass the digital assistant altogether and manually send a text message to a contact named Philippe. The digital assistant can detect that the user

took this action, infer that there was problem in its recognition of the name "Philippe," and engage the user in a dialogue to acquire a correct pronunciation.

[0011] If the digital assistant correctly recognizes a word, though, it may still mispronounce the name when it uses the name in a synthesized speech output. When this occurs, a user can indicate that the name was mispronounced, such as by saying "you said that wrong," prompting the digital assistant to engage in a dialogue to acquire a correct pronunciation.

[0012] The implementations disclosed herein provide methods, systems, computer readable storage medium and user interfaces for a digital assistant to correctly and conveniently acquire correct word pronunciations for speech recognition and synthesis.

[0013] According to some implementations, a method is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. A first speech input including at least one word is received. A first phonetic representation of the at least one word is determined, the first phonetic representation comprising a first set of phonemes selected from a speech recognition phonetic alphabet. The first set of phonemes is mapped to a second set of phonemes to generate a second phonetic representation, where the second set of phonemes are selected from a speech synthesis phonetic alphabet. The second phonetic representation is stored in association with a text string corresponding to the at least one word.

[0014] In some implementations, the text string is a name in a contact list associated with a user. In some implementations, the text string is input by a user via a keyboard. In some implementations, the text string is from a webpage displayed by the electronic device.

[0015] In some implementations, after updating the speech recognizer, a second speech input including the at least one word is received. A third phonetic representation of the at least one word is determined. It is determined that the at least one word corresponds to the text string based on a determination that the third phonetic representation is substantially similar to the first phonetic representation.

[0016] According to some implementations, a method is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. An error is detected in a speech based interaction with a digital assistant. In response to detecting the error, a speech input is received from a user, the speech input including a pronunciation of one or more words. The pronunciation is stored in association with a text string corresponding to the one or more words.

[0017] According to some implementations, a method is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. A user-specified pronunciation of a first user's name is received from the first user. The pronunciation of the first user's name is stored in association with the first user's name and a unique identifier of the first user. It is detected that a second user has created a contact item including the first user's name and the unique identifier of the first user. The user-specified pronunciation is used for one or both of recognizing the first user's name in speech inputs by the second user, and synthesizing the first user's name in speech outputs to the second user by the digital assistant.

[0018] In accordance with some implementations, an electronic device includes a display, a touch-sensitive surface, optionally one or more sensors to detect intensity of contacts with the touch-sensitive surface, one or more processors, memory, and one or more programs; the one or more programs are stored in the memory and configured to be executed by the one or more processors and the one or more programs include instructions for performing the operations of any of the methods described above. In accordance with some implementations, a computer readable storage medium has stored therein instructions which when executed by an electronic device with a display, a touch-sensitive surface, and optionally one or more sensors to detect intensity of contacts with the touch-sensitive surface, cause the device to perform the operations of any of the methods referred described above. In accordance with some implementations, an electronic device includes: a display, a touch-sensitive surface, and optionally one or more sensors to detect intensity of contacts with the touch-sensitive surface; and means for performing the operations of any of the methods described above. In accordance with some implementations, an information processing apparatus, for use in an electronic device with a display and a touch-sensitive surface, optionally one or more sensors to detect intensity of contacts with the touch-sensitive surface, includes means for performing the operations of any of the methods described above.

[0019] The details of one or more implementations of the subject matter described in this specification are set forth in the accompanying drawings and the description below. Other features, aspects, and advantages of the subject matter will become apparent from the description, the drawings, and the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0020] Figure 1 is a block diagram illustrating an environment in which a digital assistant operates in accordance with some implementations.

[0021] Figure 2 is a block diagram illustrating a digital assistant client system in accordance with some implementations.

[0022] Figure 3A is a block diagram illustrating a digital assistant system or a server portion thereof in accordance with some implementations.

[0023] Figure 3B is a block diagram illustrating functions of the digital assistant shown in Figure 3A in accordance with some implementations.

[0024] Figure 3C is a diagram of a portion of an ontology in accordance with some implementations.

[0025] Figure 4 illustrates a portion of a contact list item in accordance with some implementations.

[0026] Figures 5A-7 are flow diagrams of an exemplary method implemented by a digital assistant for learning word pronunciations, in accordance with some implementations.

[0027] Like reference numerals refer to corresponding parts throughout the drawings.

DESCRIPTION OF IMPLEMENTATIONS

[0028] Figure 1 is a block diagram of an operating environment 100 of a digital assistant according to some implementations. 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 a inferred user intent, the system can perform one or more of the following: identifying a task flow with steps and parameters designed to accomplish the inferred user intent, inputting specific requirements from the inferred user intent into the task flow; executing the task flow by invoking programs, methods, services, APIs, or the Like; and generating output responses to the user in an audible (e.g. speech) and/or visual form.

[0029] Specifically, a digital assistant is capable of accepting a user request at least partially in the form of a natural language command, request, statement, narrative, and/or inquiry. Typically, the user request seeks either an informational answer or performance of a task by the digital assistant. A satisfactory response to the user request is either provision of

the requested informational answer, performance of the requested task, or a combination of the two. For example, a user may ask the digital assistant a question, such as "Where am I right now?" Based on the user's current location, the digital assistant may answer, "You are in Central Park near the west gate." The user may also request the performance of a task, for example, "Please invite my friends to my girlfriend's birthday party next week." In response, the digital assistant may 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.

[0030] An example of a digital assistant is described in Applicant's U.S. Utility Application Serial No. 12/987,982 for "Intelligent Automated Assistant," filed January 10, 2011, the entire disclosure of which is incorporated herein by reference.

[0031] As shown in Figure 1, in some implementations, a digital assistant is implemented according to a client-server model. The digital assistant includes a client-side portion 102a, 102b (hereafter "DA client 102") executed on a user device 104a, 104b, and a server-side portion 106 (hereafter "DA server 106") executed on a server system 108. The DA client 102 communicates with the DA server 106 through one or more networks 110. The DA client 102 provides client-side functionalities such as user-facing input and output processing and communications with the DA-server 106. The DA server 106 provides server-side functionalities for any number of DA-clients 102 each residing on a respective user device 104.

[0032] In some implementations, the DA server 106 includes a client-facing I/O interface 112, one or more processing modules 114, data and models 116, and an I/O interface to external services 118. The client-facing I/O interface facilitates the client-facing input and output processing for the digital assistant server 106. The one or more processing modules 114 utilize the data and models 116 to determine the user's intent based on natural language input and perform task execution based on inferred user intent. In some implementations, the DA-server 106 communicates with external services 120 through the network(s) 110 for task

completion or information acquisition. The I/O interface to external services 118 facilitates such communications.

[0033] Examples of the user device 104 include, but are not limited to, a handheld computer, a personal digital assistant (PDA), a tablet computer, a laptop computer, a desktop computer, a cellular telephone, a smart phone, an enhanced general packet radio service (EGPRS) mobile phone, a media player, a navigation device, a game console, a television, a remote control, or a combination of any two or more of these data processing devices or other data processing devices. More details on the user device 104 are provided in reference to an exemplary user device 104 shown in Figure 2.

[0034] Examples of the communication network(s) 110 include local area networks ("LAN") and wide area networks ("WAN"), e.g., the Internet. The communication network(s) 110 may be implemented using any known network protocol, including various wired or wireless protocols, such as e.g., 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), Wi-MAX, or any other suitable communication protocol.

[0035] The server system 108 is implemented on one or more standalone data processing apparatus or a distributed network of computers. In some implementations, the server system 108 also employs various virtual devices and/or services of third party service providers (e.g., third-party cloud service providers) to provide the underlying computing resources and/or infrastructure resources of the server system 108.

[0036] Although the digital assistant shown in Figure 1 includes both a client-side portion (e.g., the DA-client 102) and a server-side portion (e.g., the DA-server 106), in some implementations, the functions of a digital assistant is implemented as a standalone application installed on a user device. In addition, the divisions of functionalities between the client and server portions of the digital assistant can vary in different implementations. For example, in some implementations, 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.

[0037] Figure 2 is a block diagram of a user-device 104 in accordance with some implementations. The user device 104 includes a memory interface 202, one or more processors 204, and a peripherals interface 206. The various components in the user device

104 are coupled by one or more communication buses or signal lines. The user device 104 includes various sensors, subsystems, and peripheral devices that are coupled to the peripherals interface 206. The sensors, subsystems, and peripheral devices gather information and/or facilitate various functionalities of the user device 104.

[0038] For example, a motion sensor 210, a light sensor 212, and a proximity sensor 214 are coupled to the peripherals interface 206 to facilitate orientation, light, and proximity sensing functions. One or more other sensors 216, such as a positioning system (e.g., GPS receiver), a temperature sensor, a biometric sensor, a gyro, a compass, an accelerometer, and the like, are also connected to the peripherals interface 206, to facilitate related functionalities.

[0039] In some implementations, a camera subsystem 220 and an optical sensor 222 are utilized to facilitate camera functions, such as taking photographs and recording video clips. Communication functions are facilitated through one or more wired and/or wireless communication subsystems 224, which can include various communication ports, radio frequency receivers and transmitters, and/or optical (e.g., infrared) receivers and transmitters. An audio subsystem 226 is coupled to speakers 228 and a microphone 230 to facilitate voice-enabled functions, such as voice recognition, voice replication, digital recording, and telephony functions.

[0040] In some implementations, an I/O subsystem 240 is also coupled to the peripheral interface 206. The I/O subsystem 240 includes a touch screen controller 242 and/or other input controller(s) 244. The touch-screen controller 242 is coupled to a touch screen 246. The touch screen 246 and the touch screen controller 242 can, for example, detect contact and movement or break thereof using any of a plurality of touch sensitivity technologies, such as capacitive, resistive, infrared, surface acoustic wave technologies, proximity sensor arrays, and the like. The other input controller(s) 244 can be coupled to other input/control devices 248, such as one or more buttons, rocker switches, thumb-wheel, infrared port, USB port, and/or a pointer device such as a stylus.

[0041] In some implementations, the memory interface 202 is coupled to memory 250. The memory 250 can include high-speed random access memory and/or non-volatile memory, such as one or more magnetic disk storage devices, one or more optical storage devices, and/or flash memory (e.g., NAND, NOR).

[0042] In some implementations, the memory 250 stores an operating system 252, a communication module 254, a user interface module 256, a sensor processing module 258, a

phone module 260, and applications 262. The operating system 252 includes instructions for handling basic system services and for performing hardware dependent tasks. The communication module 254 facilitates communicating with one or more additional devices, one or more computers and/or one or more servers. The user interface module 256 facilitates graphic user interface processing and output processing using other output channels (e.g., speakers). The sensor processing module 258 facilitates sensor-related processing and functions. The phone module 260 facilitates phone-related processes and functions. The application module 262 facilitates various functionalities of user applications, such as electronic-messaging, web browsing, media processing, Navigation, imaging and/or other processes and functions.

[0043] As described in this specification, the memory 250 also stores client-side digital assistant instructions (e.g., in a digital assistant client module 264) and various user data 266 (e.g., user-specific vocabulary data, preference data, and/or other data such as the user's electronic address book, to-do lists, shopping lists, user-specified name pronunciations, etc.) to provide the client-side functionalities of the digital assistant.

[0044] In various implementations, the digital assistant client module 264 is capable of accepting voice input (e.g., speech input), text input, touch input, and/or gestural input through various user interfaces (e.g., the I/O subsystem 244) of the user device 104. The digital assistant client module 264 is also capable of providing output in audio (e.g., speech output), visual, and/or tactile forms. For example, output can be provided as voice, sound, alerts, text messages, menus, graphics, videos, animations, vibrations, and/or combinations of two or more of the above. During operation, the digital assistant client module 264 communicates with the digital assistant server using the communication subsystems 224.

[0045] In some implementations, the digital assistant client module 264 includes a speech synthesis module 265. The speech synthesis module 265 synthesizes speech outputs for presentation to the user. The speech synthesis module 265 synthesizes speech outputs based on text provided by the digital assistant. For example, the digital assistant generates text to provide as an output to a user, and the speech synthesis module 265 converts the text to an audible speech output. The speech synthesis module 265 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.

[0046] In some implementations, the speech synthesis module 265 stores canonical pronunciations for certain words. For example, the speech synthesis module 265 may store a sequence of phonemes for the word "the" that is known to produce a correct pronunciation when synthesized. In some implementations, multiple possible pronunciations are stored for a given word, including user-specified pronunciations. As described herein, the pronunciation that is ultimately selected for synthesis is determined based on any of several possible factors or combinations thereof (e.g., the most common user-specified pronunciation, the most common user-specified pronunciation in a geographical area, etc.).

[0047] In some implementations, where a user has provided a correct or preferred pronunciation for a word (e.g., a proper name), the speech synthesis module 265 uses the user-specified pronunciation for those words to the exclusion of other possible pronunciations. (Techniques for acquiring and processing user-specified pronunciations are discussed herein.) In some implementations, user-specified pronunciations for use by the speech synthesis module 265 are represented using a speech synthesis phonetic alphabet (e.g., an alphabet or other symbolic linguistic representation used by the speech synthesis module 265 to synthesize speech outputs).

[0048] In some implementations, the user-specified pronunciations are stored in the user data 266. For example, user-specified pronunciations of the names of contacts in a user's electronic address book or contact list are stored in association with the respective contacts. User-specified pronunciations may be visible or hidden to the user. When they are visible and/or accessible to the user, they may be manually correctable through the electronic address book. For example, a user can select a user-specified pronunciation and modify, alter, or replace it, using text or speech inputs.

[0049] In some implementations, user-specified pronunciations of other words (e.g., words that are not names of a user's contacts) are stored in user-specific vocabularies in the user data 266. Thus, in some implementations, any words for which the user wishes to specify a particular pronunciation are accessible by the speech synthesis module 265.

[0050] In some implementations, user-specified pronunciations are stored remotely from the user device 104, such as in a remote server or cloud-based service (e.g., server system 108, Figure 1). In such cases, the user-specified pronunciations are still associated with the user, and may be encrypted or otherwise secured so that only an authorized user and/or the authorized user's devices can access the information. Accordingly, user-specified pronunciations of words are accessible to a user via multiple user devices. This also helps

increase the perceived intelligence of the digital assistant, because once a user specifies a particular pronunciation of a word or name, the digital assistant can use the correct pronunciation regardless of whether the user is interacting with the digital assistant on her smart phone or other computing device, e.g., laptop computer or tablet.

[0051] In some implementations, user-specified pronunciations are stored both locally (e.g., on one or more user devices 104) and remotely (e.g., on the server system 108). In some implementations, user-specified pronunciations for a particular user are copied to a user device upon authentication of the device to access an account associated with the user. For example, user-specified pronunciations stored on the server system 108 may be associated with a particular user account, and when a device becomes associated with that user account (e.g., because the user logged into his or her account on that device), user data (e.g., user data 266) for that account is sent to or otherwise becomes accessible by the device.

[0052] In some implementations, instead of (or in addition to) using the local speech synthesis module 265, speech synthesis is performed on a remote device (e.g., the server system 108), and the synthesized speech is sent to the user device 104 for output to the user. For example, this occurs 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 may be possible to obtain higher quality speech outputs than would be practical with client-side synthesis.

[0053] In some implementations, the digital assistant client module 264 utilizes the various sensors, subsystems and peripheral devices to gather additional information from the surrounding environment of the user device 104 to establish a context associated with a user, the current user interaction, and/or the current user input. In some implementations, the digital assistant client module 264 provides the context information or a subset thereof with the user input to the digital assistant server to help infer the user's intent. In some implementations, the digital assistant also uses the context information to determine how to prepare and delivery outputs to the user.

[0054] In some implementations, the context information that accompanies the user input includes sensor information, e.g., lighting, ambient noise, ambient temperature, images or videos of the surrounding environment, etc. In some implementations, the context information also includes 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 implementations, information related to the software state of

the user device 106, e.g., running processes, installed programs, past and present network activities, background services, error logs, resources usage, etc., of the user device 104 are provided to the digital assistant server as context information associated with a user input.

[0055] In some implementations, the DA client module 264 selectively provides information (e.g., user data 266) stored on the user device 104 in response to requests from the digital assistant server. In some implementations, the digital assistant client module 264 also elicits additional input from the user via a natural language dialogue or other user interfaces upon request by the digital assistant server 106. The digital assistant client module 264 passes the additional input to the digital assistant server 106 to help the digital assistant server 106 in intent deduction and/or fulfillment of the user's intent expressed in the user request.

[0056] In various implementations, the memory 250 includes additional instructions or fewer instructions. Furthermore, various functions of the user device 104 may be implemented in hardware and/or in firmware, including in one or more signal processing and/or application specific integrated circuits.

[0057] Figure 3A is a block diagram of an example digital assistant system 300 in accordance with some implementations. In some implementations, the digital assistant system 300 is implemented on a standalone computer system. In some implementations, the digital assistant system 300 is distributed across multiple computers. In some implementations, some of the modules and functions of the digital assistant are divided into a server portion and a client portion, where the client portion resides on a user device (e.g., the user device 104) and communicates with the server portion (e.g., the server system 108) through one or more networks, e.g., as shown in Figure 1. In some implementations, the digital assistant system 300 is an implementation of the server system 108 (and/or the digital assistant server 106) shown in Figure 1. It should be noted that the digital assistant system 300 is only one example of a digital assistant system, and that the digital assistant system 300 may have more or fewer components than shown, may combine two or more components, or may have a different configuration or arrangement of the components. The various components shown in Figure 3A may be implemented in hardware, software instructions for execution by one or more processors, firmware, including one or more signal processing and/or application specific integrated circuits, or a combination of thereof.

[0058] The digital assistant system 300 includes memory 302, one or more processors 304, an input/output (I/O) interface 306, and a network communications interface 308. These

components communicate with one another over one or more communication buses or signal lines 310.

[0059] In some implementations, the memory 302 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).

[0060] In some implementations, the I/O interface 306 couples input/output devices 316 of the digital assistant system 300, such as displays, keyboards, touch screens, and microphones, to the user interface module 322. The I/O interface 306, in conjunction with the user interface module 322, receives user inputs (e.g., voice input, keyboard inputs, touch inputs, etc.) and processes them accordingly. In some implementations, e.g., when the digital assistant is implemented on a standalone user device, the digital assistant system 300 includes any of the components and I/O and communication interfaces described with respect to the user device 104 in Figure 2. In some implementations, the digital assistant system 300 represents the server portion of a digital assistant implementation, and interacts with the user through a client-side portion residing on a user device (e.g., the user device 104 shown in Figure 2).

[0061] In some implementations, the network communications interface 308 includes wired communication port(s) 312 and/or wireless transmission and reception circuitry 314. The wired communication port(s) receive and send communication signals via one or more wired interfaces, e.g., Ethernet, Universal Serial Bus (USB), FIREWIRE, etc. The wireless circuitry 314 receives and sends RF signals and/or optical signals from/to communications networks and other communications devices. The wireless communications may use any of a plurality of communications standards, protocols and technologies, such as GSM, EDGE, CDMA, TDMA, Bluetooth, Wi-Fi, VoIP, Wi-MAX, or any other suitable communication protocol. The network communications interface 308 enables communication between the digital assistant system 300 with networks, such as the Internet, an intranet and/or a wireless network, such as a cellular telephone network, a wireless local area network (LAN) and/or a metropolitan area network (MAN), and other devices.

[0062] In some implementations, memory 302, or the computer readable storage media of memory 302, stores programs, modules, instructions, and data structures including all or a subset of: an operating system 318, a communications module 320, a user interface module 322, one or more applications 324, and a digital assistant module 326. The one or more

processors 304 execute these programs, modules, and instructions, and reads/writes from/to the data structures.

[0063] The operating system 318 (e.g., Darwin, RTXC, LINUX, UNIX, OS X, WINDOWS, or an embedded operating system such as VxWorks) includes various software components and/or drivers for controlling and managing general system tasks (e.g., memory management, storage device control, power management, etc.) and facilitates communications between various hardware, firmware, and software components.

[0064] The communications module 320 facilitates communications between the digital assistant system 300 with other devices over the network communications interface 308. For example, the communication module 320 may communicate with the communication interface 254 of the device 104 shown in Figure 2. The communications module 320 also includes various components for handling data received by the wireless circuitry 314 and/or wired communications port 312.

[0065] The user interface module 322 receives commands and/or inputs from a user via the I/O interface 306 (e.g., from a keyboard, touch screen, pointing device, controller, and/or microphone), and generates user interface objects on a display. The user interface module 322 also prepares and delivers outputs (e.g., speech, sound, animation, text, icons, vibrations, haptic feedback, and light, etc.) to the user via the I/O interface 306 (e.g., through displays, audio channels, speakers, and touch-pads, etc.).

[0066] The applications 324 include programs and/or modules that are configured to be executed by the one or more processors 304. For example, if the digital assistant system is implemented on a standalone user device, the applications 324 may include user applications, such as games, a calendar application, a navigation application, or an email application. If the digital assistant system 300 is implemented on a server farm, the applications 324 may include resource management applications, diagnostic applications, or scheduling applications, for example.

[0067] The memory 302 also stores the digital assistant module (or the server portion of a digital assistant) 326. In some implementations, the digital assistant module 326 includes the following sub-modules, or a subset or superset thereof: an input/output processing module 328, a speech-to-text (STT) processing module 330, a phonetic alphabet conversion module 331, a natural language processing module 332, a dialogue flow processing module 334, a task flow processing module 336, a service processing module 338, and a speech interaction error detection module 339. Each of these modules has access to one or more of the

following data and models of the digital assistant 326, or a subset or superset thereof: ontology 360, vocabulary index 344, user data 348, task flow models 354, and service models 356.

[0068] In some implementations, using the processing modules, data, and models implemented in the digital assistant module 326, the digital assistant performs at least some of the following: identifying a user's intent expressed in a natural language input received from the user; actively eliciting and obtaining information needed to fully infer the user's intent (e.g., by disambiguating words, names, intentions, etc.); determining the task flow for fulfilling the inferred intent; and executing the task flow to fulfill the inferred intent.

[0069] In some implementations, as shown in Figure 3B, the I/O processing module 328 interacts with the user through the I/O devices 316 in Figure 3A or with a user device (e.g., a user device 104 in Figure 1) through the network communications interface 308 in Figure 3A to obtain user input (e.g., a speech input) and to provide responses (e.g., as speech outputs) to the user input. The I/O processing module 328 optionally obtains context information associated with the user input from the user device, along with or shortly after the receipt of the user input. The context information includes user-specific data, vocabulary, and/or preferences relevant to the user input. In some implementations, the context information also includes software and hardware states of the device (e.g., the user device 104 in Figure 1) at the time the user request is received, and/or information related to the surrounding environment of the user at the time that the user request was received. In some implementations, the I/O processing module 328 also sends follow-up questions to, and receives answers from, the user regarding the user request. When a user request is received by the I/O processing module 328 and the user request contains a speech input, the I/O processing module 328 forwards the speech input to the speech-to-text (STT) processing module 330 for speech-to-text conversions.

[0070] The speech-to-text processing module 330 (or speech recognizer) receives speech input (e.g., a user utterance captured in a voice recording) through the I/O processing module 328. In some implementations, the STT processing module 330 uses various acoustic and language models to recognize the speech input as a sequence of phonemes, and ultimately, a sequence of words or tokens written in one or more languages. The STT processing module 330 can be implemented using any suitable speech recognition techniques, acoustic models, and language models, such as Hidden Markov Models, Dynamic Time Warping (DTW)-based speech recognition, and other statistical and/or analytical techniques. In some

implementations, the speech-to-text processing can be performed at least partially by a third party service or on the user's device. Once the STT processing module 330 obtains the result of the speech-to-text processing, e.g., a sequence of words or tokens, it passes the result to the natural language processing module 332 for intent deduction. In some implementations, the STT module 330 resides on a server computer (e.g., the server system 108), while in some implementations, it resides on a client device (e.g., the user device 104).

[0071] In some implementations, the STT processing module 330 includes and/or accesses a vocabulary of recognizable words. Each word is associated with one or more candidate pronunciations of the word represented in a speech recognition phonetic alphabet. For example, the vocabulary may include the word "tomato" in association with the candidate pronunciations of "tuh-may-doe" and "tuh-mah-doe." In some implementations, 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 implementations, the candidate pronunciations are manually generated, e.g., based on known canonical pronunciations.

[0072] In some implementations, the candidate pronunciations are ranked based on the commonness of the candidate pronunciation. For example, the candidate pronunciation "tuh-may-doe" may be ranked higher than "tuh-mah-doe," because the former is a more commonly used pronunciation in the user's geographic region. In some implementations, one of the candidate pronunciations is selected as a predicted pronunciation (e.g., the most likely pronunciation).

[0073] When an utterance is received, the STT processing module 330 attempts to identify the phonemes in the utterance (e.g., using an acoustic model), and then attempts to identify words that match the phonemes (e.g., using a language model). For example, if the STT processing module 330 first identifies the sequence of phonemes "tuh-may-doe" in an utterance, it then determines, based on the vocabulary 344, that this sequence corresponds to the word "tomato."

[0074] In some implementations, the STT processing module 330 uses approximate matching techniques to determine words in an utterance. Thus, for example, the STT processing module 330 can determine that the sequence of phonemes "duh-may-doe" corresponds to the word "tomato," even if that particular sequence of phonemes is not one of the candidate pronunciations for that word.

[0075] As described herein, in some implementations, the STT processing module 330 identifies phonemes in an utterance of a known word for the purpose of generating a user-

specified pronunciation of the word. Thus, for example, a user can add a user-specified pronunciation for the word "tomato" by simply speaking the preferred pronunciation to the digital assistant. The STT processing module 330 processes the utterance containing the preferred pronunciation to identify the phonemes in the utterance. For example, the STT processing module 330 may identify the phonemes "tuh-may-duh" in the utterance, and store that set of phonemes (e.g., in the vocabulary index 344 and/or user data 348) as a user-specified pronunciation for the word "tomato."

[0076] In some implementations, the speech-to-text processing module 330 uses user-specified pronunciations to help recognize certain words in user utterances. For example, a user may discover that the digital assistant cannot accurately recognize a particular contact's name. By specifying a preferred pronunciation for the name, the digital assistant, and specifically the speech-to-text processing module 330, will thereafter accurately recognize the name in user utterances.

[0077] In some implementations, user-specified pronunciations for speech recognition by the speech-to-text processing module 330 are stored in the vocabulary index 344. In some implementations, user-specified pronunciations are also or instead stored in association with words in user data 348. For example, if a user specifies a pronunciation for a name in her contact list (which is stored in user data 348), the user-specified pronunciation is stored in association with the contact and/or the name in user data 348 or 266. In some implementations, user-specified name pronunciations are visible to the user, while in implementations they are not.

[0078] In some implementations, all user-specified pronunciations are stored in user data 348 or 266, and are accessed by the STT processing module 330 and/or incorporated into the vocabulary index 344 when appropriate. Thus, user data for specific users is stored in association with a user account, and is used to augment and/or customize a generic STT processing module and/or vocabulary index, for example, when a user authorizes a new device to access his account.

[0079] In some implementations, user-specified pronunciations for use by the speech-to-text processing module 330 are represented using a speech recognition phonetic alphabet (e.g., an alphabet or other symbolic linguistic representation used by the speech-to-text processing module 330 to recognize speech inputs). In some implementations, the speech recognition phonetic alphabet corresponds to the set of phonemes that the STT processing module 330 is capable of identifying in a recording of a spoken utterance. In some

implementations, the speech recognition alphabet is the International Phonetic Alphabet ("IPA"), or a subset of the IPA that contains phonemes likely to be used by users in a particular geographical region or by speakers of a particular language.

[0080] The phonetic alphabet conversion module 331 converts phonetic representations of words between different phonetic alphabets. Specifically, in some implementations, a speech recognizer (e.g., the STT processing module 330) uses a speech recognition phonetic alphabet to determine the phonemes in an utterance, while a speech synthesizer (e.g., the speech synthesis module 265, Figure 2) uses a speech synthesis phonetic alphabet that is different from the speech recognition phonetic alphabet to synthesize a speech output. Speech synthesizers and speech recognizers, therefore, cannot share a single phonetic representation because they use different phonetic alphabets. Thus, in some implementations, the phonetic alphabet conversion module 331 converts phonetic representations from one phonetic alphabet (e.g., a speech recognition phonetic alphabet) into a second phonetic alphabet that is different than the first (e.g., a speech synthesis phonetic alphabet). Accordingly, as described herein, a phonetic representation of a word that is determined using the STT processing module 330 can be converted or mapped to a phonetic alphabet that is usable by the speech synthesis module 265.

[0081] The natural language processing module 332 ("natural language processor") of the digital assistant takes the sequence of words or tokens ("token sequence") generated by the speech-to-text processing module 330, and attempts to associate the token sequence with one or more "actionable intents" recognized by the digital assistant. An "actionable intent" represents a task that can be performed by the digital assistant, and has an associated task flow implemented in the task flow models 354. The associated task flow is a series of programmed actions and steps that the digital assistant takes in order to perform the task. The scope of a digital assistant's capabilities is dependent on the number and variety of task flows that have been implemented and stored in the task flow models 354, or in other words, on the number and variety of "actionable intents" that the digital assistant recognizes. The effectiveness of the digital assistant, however, is also dependent on the assistant's ability to infer the correct "actionable intent(s)" from the user request expressed in natural language.

[0082] In some implementations, in addition to the sequence of words or tokens obtained from the speech-to-text processing module 330, the natural language processor 332 also receives context information associated with the user request, e.g., from the I/O processing module 328. The natural language processor 332 optionally uses the context information to

clarify, supplement, and/or further define the information contained in the token sequence received from the speech-to-text processing module 330. The context information includes, for example, user preferences, hardware and/or software states of the user device, sensor information collected before, during, or shortly after the user request, prior interactions (e.g., dialogue) between the digital assistant and the user, and the like. As described in this specification, context information is dynamic, and can change with time, location, content of the dialogue, and other factors.

[0083] In some implementations, the natural language processing is based on e.g., ontology 360. The ontology 360 is a hierarchical structure containing many nodes, each node representing either an "actionable intent" or a "property" relevant to one or more of the "actionable intents" or other "properties". As noted above, an "actionable intent" represents a task that the digital assistant is capable of performing, i.e., it is "actionable" or can be acted on. A "property" represents a parameter associated with an actionable intent or a sub-aspect of another property. A linkage between an actionable intent node and a property node in the ontology 360 defines how a parameter represented by the property node pertains to the task represented by the actionable intent node.

[0084] In some implementations, the ontology 360 is made up of actionable intent nodes and property nodes. Within the ontology 360, each actionable intent node is linked to one or more property nodes either directly or through one or more intermediate property nodes. Similarly, each property node is linked to one or more actionable intent nodes either directly or through one or more intermediate property nodes. For example, as shown in Figure 3C, the ontology 360 may include 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). 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 Figure 3C, the ontology 360 may also include a "set reminder" node (i.e., another actionable intent node). Property nodes "date/time" (for the setting the reminder) and "subject" (for the reminder) are each linked to the "set reminder" node. Since the property "date/time" is relevant to both the task of making a restaurant reservation and the task of setting a reminder, the property node "date/time" is

linked to both the "restaurant reservation" node and the "set reminder" node in the ontology 360.

[0085] An actionable intent node, along with its linked concept nodes, may be described as a "domain." In the present discussion, each domain is associated with a respective actionable intent, and refers to the group of nodes (and the relationships therebetween) associated with the particular actionable intent. For example, the ontology 360 shown in Figure 3C includes an example of a restaurant reservation domain 362 and an example of a reminder domain 364 within the ontology 360. The restaurant reservation domain includes the actionable intent node "restaurant reservation," property nodes "restaurant," "date/time," and "party size," and sub-property nodes "cuisine," "price range," "phone number," and "location." The reminder domain 364 includes the actionable intent node "set reminder," and property nodes "subject" and "date/time." In some implementations, the ontology 360 is made up of many domains. Each domain may share one or more property nodes with one or more other domains. For example, the "date/time" property node may be associated with many different domains (e.g., a scheduling domain, a travel reservation domain, a movie ticket domain, etc.), in addition to the restaurant reservation domain 362 and the reminder domain 364.

[0086] While Figure 3C illustrates two example domains within the ontology 360, other domains (or actionable intents) include, for example, "initiate a phone call," "find directions," "schedule a meeting," "send a message," and "provide an answer to a question," "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 may further include property nodes such as "recipient(s)", "message type", and "message body." The property node "recipient" may be further defined, for example, by the sub-property nodes such as "recipient name" and "message address."

[0087] In some implementations, the ontology 360 includes all the domains (and hence actionable intents) that the digital assistant is capable of understanding and acting upon. In some implementations, the ontology 360 may be modified, such as by adding or removing entire domains or nodes, or by modifying relationships between the nodes within the ontology 360.

[0088] In some implementations, nodes associated with multiple related actionable intents may be clustered under a "super domain" in the ontology 360. For example, a "travel" super-domain may include a cluster of property nodes and actionable intent nodes related to travels. The actionable intent nodes related to travels may include "airline reservation," "hotel

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

[0089] In some implementations, each node in the ontology 360 is associated with a set of words and/or phrases that are relevant to the property or actionable intent represented by the node. The respective set of words and/or phrases associated with each node is the so-called "vocabulary" associated with the node. The respective set of words and/or phrases associated with each node can be stored in the vocabulary index 344 in association with the property or actionable intent represented by the node. For example, returning to Figure 3B, the vocabulary associated with the node for the property of "restaurant" may include 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" may include words and phrases such as "call," "phone," "dial," "ring," "call this number," "make a call to," and so on. The vocabulary index 344 optionally includes words and phrases in different languages.

[0090] The natural language processor 332 receives the token sequence (e.g., a text string) from the speech-to-text processing module 330, and determines what nodes are implicated by the words in the token sequence. In some implementations, if a word or phrase in the token sequence is found to be associated with one or more nodes in the ontology 360 (via the vocabulary index 344), the word or phrase will "trigger" or "activate" those nodes. Based on the quantity and/or relative importance of the activated nodes, the natural language processor 332 will select one of the actionable intents as the task that the user intended the digital assistant to perform. In some implementations, the domain that has the most "triggered" nodes is selected. In some implementations, the domain having the highest confidence value (e.g., based on the relative importance of its various triggered nodes) is selected. In some implementations, the domain is selected based on a combination of the number and the importance of the triggered nodes. In some implementations, 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.

[0091] In some implementations, the digital assistant also stores names of specific entities in the vocabulary index 344, so that when one of these names is detected in the user request, the natural language processor 332 will be able to recognize that the name refers to a specific instance of a property or sub-property in the ontology. In some implementations, the names of specific entities are names of businesses, restaurants, people, movies, and the like. In some implementations, the digital assistant searches and identifies specific entity names from other data sources, such as the user's address book, a movies database, a musicians database, and/or a restaurant database. In some implementations, when the natural language processor 332 identifies that a word in the token sequence *is* a name of a specific entity (such as a name in the user's address book), that word is given additional significance in selecting the actionable intent with *in* the ontology for the user request.

[0092] For example, when the words "Mr. Santo" are recognized from the user request, and the last name "Santo" is found in the vocabulary index 344 as one of the contacts in the user's contact list, then it is likely that the user request corresponds to a "send a message" or "initiate a phone call domain. For another example, when the words "ABC Cafe" are found in the user request, and the term "ABC Cafe" is found in the vocabulary index 344 as the name of a particular restaurant in the user's city, then it is likely that the user request corresponds to a "restaurant reservation" domain.

[0093] User data 348 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 implementations, the natural language processor 332 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," the natural language processor 332 is able to access user data 348 to determine who the "friends" are and when and where the "birthday party" would be held, rather than requiring the user to provide such information explicitly in his/her request.

[0094] Other details of searching an ontology based on a token string is described in U.S. Utility Application Serial No. 12/341,743 for "Method and Apparatus for Searching Using An Active Ontology," filed December 22, 2008, the entire disclosure of which is incorporated herein by reference.

[0095] In some implementations, once the natural language processor 332 identifies an actionable intent (or domain) based on the user request, the natural language processor 332

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

[0096] In some implementations, the natural language processor 332 passes the structured query (including any completed parameters) to the task flow processing module 336 ("task flow processor"). The task flow processor 336 is configured to receive the structured query from the natural language processor 332, complete the structured query, if necessary, and perform the actions required to "complete" the user's ultimate request. In some implementations, the various procedures necessary to complete these tasks are provided in task flow models 354. In some implementations, the task flow models include procedures for obtaining additional information from the user, and task flows for performing actions associated with the actionable intent.

[0097] As described above, in order to complete a structured query, the task flow processor 336 may need to initiate additional dialogue with the user in order to obtain additional information, and/or disambiguate potentially ambiguous utterances. When such interactions are necessary, the task flow processor 336 invokes the dialogue processing module 334

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

[0098] In some cases, the task flow processor 336 may receive a structured query that has one or more ambiguous properties. For example, a structured query for the "send a message" domain may indicate that the intended recipient is "Bob," and the user may have multiple contacts named "Bob." The task flow processor 336 will request that the dialogue processor 334 disambiguate this property of the structured query. In turn, the dialogue processor 334 may ask the user "Which Bob?", and display (or read) a list of contacts named "Bob" from which the user may choose.

[0099] Once the task flow processor 336 has completed the structured query for an actionable intent, the task flow processor 336 proceeds to perform the ultimate task associated with the actionable intent. Accordingly, the task flow processor 336 executes the steps and instructions in the task flow model according to the specific parameters contained in the structured query. For example, the task flow model for the actionable intent of "restaurant reservation" may include steps and instructions for contacting a restaurant and actually requesting a reservation for a particular party size at a particular time. For example, using a structured query such as: {restaurant reservation, restaurant = ABC Cafe, date = 3/12/2012, time = 7pm, party size = 5}, the task flow processor 336 may perform the steps of: (1) logging onto a server of the ABC Cafe or a restaurant reservation system such as OPENTABLE®, (2) entering the date, time, and party size information in a form on the website, (3) submitting the form, and (4) making a calendar entry for the reservation in the user's calendar.

[0100] In some implementations, the task flow processor 336 employs the assistance of a service processing module 338 ("service processor") to complete a task requested in the user input or to provide an informational answer requested in the user input. For example, the service processor 338 can act on behalf of the task flow processor 336 to make a phone call, set a calendar entry, invoke a map search, invoke or interact with other user applications installed on the user device, and invoke or interact with third party services (e.g. a restaurant reservation portal, a social networking website, a banking portal, etc.). In some implementations, the protocols and application programming interfaces (API) required by each service can be specified by a respective service model among the services models 356. The service processor 338 accesses the appropriate service model for a service and generates requests for the service in accordance with the protocols and APIs required by the service according to the service model.

[0101] For example, if a restaurant has enabled an online reservation service, the restaurant can submit 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 the task flow processor 336, the service processor 338 can establish a network connection with the online reservation service using the web address stored in the service model, and send the necessary parameters of the reservation (e.g., time, date, party size) to the online reservation interface in a format according to the API of the online reservation service.

[0102] In some implementations, the natural language processor 332, dialogue processor 334, and task flow processor 336 are used collectively and iteratively to infer and define the user's intent, obtain information to further clarify and refine the user intent, and finally generate a response (i.e., an output to the user, or the completion of a task) to fulfill the user's intent.

[0103] In some implementations, after all of the tasks needed to fulfill the user's request have been performed, the digital assistant 326 formulates a confirmation response, and sends the response back to the user through the I/O processing module 328. If the user request seeks an informational answer, the confirmation response presents the requested information to the user. In some implementations, the digital assistant also requests the user to indicate whether the user is satisfied with the response produced by the digital assistant 326.

[0104] The error detection module 339 detects errors in interactions between a user and the digital assistant. In some implementations, to detect errors, the error detection module 339

monitors interactions between a user and the digital assistant, and/or between a user and a user device. For example, the error detection module 339 monitors any of the following types of interactions, or a subset thereof: a user's speech inputs to the digital assistant (e.g., if a user says "you got that wrong" or "you are pronouncing that wrong"), button presses (e.g., if a user selects a lock-screen or "home" button (or any other affordance) to cancel an action), movements of the device (e.g., shaking the device, setting the device down in a certain orientation, such as screen-down), termination of actions or suggested actions on the user device (e.g., cancelling a telephone call, email, text message, etc. after the digital assistant initiates or suggests it), initiation of an action shortly after a digital assistant fails to successfully infer an intent or adequately respond to a user, etc. In some implementations, the error detection module 339 monitors other types of interactions to detect errors as well.

[0105] In order to detect such errors, in some implementations, the error detection module 339 communicates with or otherwise receives information from various modules and components of the digital assistant system 300 and/or the user device 104, such as the I/O processing module 328 (and/or the I/O devices 316), the STT processing module 330, natural language processing module 332, the dialogue flow processing module 334, the task flow processing module 336, the service processing module 338, the phone module 260, the sensor processing module 258, the I/O subsystem 240, and/or any of the sensors or I/O devices associated therewith.

[0106] In some implementations, the error detection module 339 monitors actions taken by the user on the user device 104 after the user cancels an action and/or dialogue with the digital assistant prior to the completion of the action or dialogue. In particular, actions taken by the user after such an occurrence often indicate both that the digital assistant did not accurately infer the user's intent (or did not understand the user's speech input), and what the digital assistant should have done based on the user's input. As a specific example, a user may ask the digital assistant to "Call Philippe," and the digital assistant may respond by saying "Calling Phil." The user may quickly cancel the telephone call to Phil, exit the dialogue with the digital assistant, and proceed to manually initiate a telephone call with a contact named Philippe. Accordingly, the error detection module 339 detects that because the telephone call to Phil was canceled, an error was made, and that the speech input to call "Philippe" should be associated with the contact that was manually selected by the user.

[0107] More details on the digital assistant can be found in the U.S. Utility Application No. 12/987,982, entitled "Intelligent Automated Assistant", filed January 10, 2011, U.S. Utility

Application No. 611493,201, entitled "Generating and Processing Data Items That Represent Tasks to Perform", filed June 3, 2011, the entire disclosures of which are incorporated herein by reference.

[0108] In most scenarios, when the digital assistant receives a user input from a user, the digital assistant attempts to provide an appropriate response to the user input with as little delay as possible. For example, suppose the user requests certain information (e.g., current traffic information) by providing a speech input (e.g., "How does the traffic Look right now?"). Right after the digital assistant receives and processes the speech input, the digital assistant optionally provides a speech output (e.g., "Looking up traffic information. . .") acknowledging receipt of the user request. After the digital assistant obtains the requested information in response to the user request, the digital assistant proceeds to provide the requested information to the user without further delay. For example, in response to the user's traffic information request, the digital assistant may provide a series of one or more discrete speech outputs separated by brief pauses (e.g., "There are 2 accidents on the road. <Pause> One accident is on 101 north bound near Whipple Avenue. <Pause> And a second accident is on 85 north near 280."), immediately after the speech outputs are generated.

[0109] For the purpose of this specification, the initial acknowledgement of the user request and the series of one or more discrete speech outputs provided in response to the user request are all considered sub-responses of a complete response to the user request. In other words, the digital assistant initiates an information provision process for the user request upon receipt of the user request, and during the information provision process, the digital assistant prepares and provides each sub-response of the complete response to the user request without requiring further prompts from the user.

[0110] Sometimes, additional information or clarification (e.g., route information) is required before the requested information can be obtained. In such scenarios, the digital assistant outputs a question (e.g., "Where are you going?") to the user asking for the additional information or clarification. In some implementations, the question provided by the digital assistant is considered a complete response to the user request because the digital assistant will not take further actions or provide any additional response to the user request until a new input is received from the user. In some implementations, once the user provides the additional information or clarification, the digital assistant initiates a new information provision process for a "new" user request established based on the original user request and the additional user input.

[0111] In some implementations, the digital assistant initiates a new information provision process upon receipt of each new user input, and each existing information provision process terminates either (1) when all of the sub-responses of a complete response to the user request have been provided to the user or (2) when the digital assistant provides a request for additional information or clarification to the user regarding a previous user request that started the existing information provision process.

[0112] In general, after a user request for information or performance of a task is received by the digital assistant, it is desirable that the digital assistant provides a response (e.g., either an output containing the requested information, an acknowledgement of a requested task, or an output to request a clarification) as promptly as possible. Real-time responsiveness of the digital assistant is one of the key factors in evaluating performance of the digital assistant. In such cases, a response is prepared as quickly as possible, and a default delivery time for the response is a time immediately after the response is prepared.

[0113] Sometimes, however, after an initial sub-response provided immediately after receipt of the user input, the digital assistant provides the remaining one or more sub-responses one at a time over an extended period of time. In some implementations, the information provision process for a user request is stretched out over an extended period of time that is longer than the sum of the time required to provide each sub-response individually. For example, in some implementations, short pauses (i.e., brief periods of silence) are inserted between an adjacent pair of sub-responses (e.g., a pair of consecutive speech outputs) when they are delivered to the user through an audio-output channel.

[0114] In some implementations, a sub-response is held in abeyance after it is prepared and is delivered only when a predetermined condition has been met. In some implementations, the predetermined condition is met when a predetermined trigger time has been reached according to a system clock and/or when a predetermined trigger event has occurred. For example, if the user says to the digital assistant "set me a timer for 5 minutes," the digital assistant initiates an information provision process upon receipt of the user request. During the information provision process, the digital assistant provides a first sub-response (e.g., "OK, timer started.") right away, and does not provide a second and final sub-response (e.g., "OK, five minutes are up") until 5 minutes later. In such cases, the default delivery time for the first sub-response is a time immediately after the first sub-response is prepared, and the default delivery time for the second, final sub-response is a time immediately after the occurrence of the trigger event (e.g., the elapse of 5 minutes from the start of the timer). The

information provision process is terminated when the digital assistant finishes providing the final sub-response to the user. In various implementations, the second sub-response is prepared any time (e.g., right after the first sub-response is prepared, or until shortly before the default delivery time for the second sub-response) before the default delivery time for the second sub-response.

[0115] Figure 4 illustrates a portion of a contact list entry 400, according to some implementations. The contact list entry 400 includes a user-specified pronunciation for a name and a standard pronunciation for the name, each stored in association with the name. While Figure 4 depicts a contact list entry, the figure and associated description applies equally to words from other lists, databases, vocabularies, etc., such as the vocabulary index 344. In some implementations, the contact list entry 400 is stored in user data 348 or 266.

[0116] The contact list entry 400 includes a name 402 ("Philippe"). The contact list entry 400 also includes a user-specified pronunciation for the name. The user-specified pronunciations for this name may be obtained and/or generated as described elsewhere in this application. Because, in some implementations, the STT processing module 330 and the speech synthesis module 265 use different phonetic alphabets, the user-specified pronunciation includes a phonetic representation of the name 402 in a speech recognition alphabet (phonetic representation 404), as well as a phonetic representation of the name 402 in a speech synthesis alphabet (phonetic representation 406). Both the representation 404 in the recognition alphabet and the representation 406 in the synthesis alphabet are based on the same pronunciation, and, therefore, the user's preferred pronunciation will both be accurately recognized by the STT processing module 330 and accurately synthesized by the speech synthesis module 265.

[0117] The phonetic representations in Figure 4 are used herein merely to illustrate that, although they represent the same pronunciation, phonetic representations for recognition and synthesis are different from one another, and that the phonemes are selected from different phonetic alphabets and/or sets of phonemes. The particular phonetic representations in Figure 4 do not necessarily reflect any particular phonetic alphabet, any particular phonemes for use by speech synthesizers or recognizers, nor any particular phoneme sequences that may be associated with the word shown. Rather, they are intended to help illustrate the differences between different pronunciations (e.g., standard and user-specified pronunciations) and between phoneme sequences for speech recognition and speech synthesis.

[0118] The contact list entry 400 also includes a standard pronunciation including a phonetic representation of the name 402 in a speech recognition alphabet (phonetic representation 408), as well as a phonetic representation of the name 402 in a speech synthesis alphabet (phonetic representation 410). The standard pronunciation represents a default pronunciation for recognition and synthesis that is used by the digital assistant absent any user-specified pronunciations for that word. In some implementations, the digital assistant maintains the standard pronunciation in addition to any user-specified pronunciations.

[0119] In some implementations, contact list entries (or any word in the vocabulary 344 and/or user data 348) include more or less information than depicted in Figure 4. Moreover, the user-specified pronunciation and the standard pronunciation (and their respective recognition and synthesis phonetic representations) need not be stored in a common location or device. For example, in some implementations, the speech recognition phonetic representations 404 and 408 are stored on a server (e.g., the server system 108) in association with a STT processing module (e.g., the STT processing module 330), while the speech synthesis phonetic representations 406 and 410 are stored on a user device (e.g., user device 104) in association with a speech synthesis module (e.g., the speech synthesis module 265). In some implementations, all of the phonetic representations are stored on both a server and a user device. In some implementations, the speech recognition phonetic representations 404 and 408 and the speech synthesis phonetic representations 406 and 410 are stored in any appropriate combination on either or both of a server and a user device.

[0120] Furthermore, in some implementations, user-specified pronunciations are stored in association with a user's contacts (or other user-specific data), while standard pronunciations are not stored in association with the user, but are part of the generic STT processing and speech synthesis modules used by multiple different instances of a digital assistant.

[0121] Figures 5A-5B are flow diagrams of an exemplary method 500 implemented by a digital assistant for learning word pronunciations. In some implementations, the method 500 is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. For example, in some implementations, the method 500 is performed at the user device 104 and/or the server system 108. In some implementations, the method 500 is performed by the digital assistant system 300 (Figure 3A), which, as noted above, may be implemented on a standalone computer system (e.g., either the user device 104 or the server system 108) or distributed

across multiple computers (e.g., the user device 104, the server system 108, and/or additional or alternative devices or systems). While the following discussion describes the method 500 as being performed by a digital assistant (e.g., the digital assistant system 300), the method is not limited to performance by any particular device or combination of devices. Moreover, the individual steps of the method may be distributed among the one or more computers, systems, or devices in any appropriate manner.

[0122] The digital assistant receives a first speech input including at least one word (502). In some implementations, the speech input corresponds to a user utterance recorded and/or received by the user device 104. In some implementations, the first input is received in the course of an interaction with the digital assistant. In some implementations, the word is a name, such as a name of a contact associated with a user. As a specific example, the utterance may be "Call Philippe Martin," corresponding to a request by the user to initiate a telephone call with a contact named Philippe Martin. In some implementations, the first input is received during an interaction in which the user is specifically training the digital assistant how to pronounce a particular word. As a specific example, the utterance may be "Please pronounce that 'Philippe,'" or, if the digital assistant has already prompted the user to provide a preferred pronunciation for a particular word, the utterance may simply be "Philippe."

[0123] The digital assistant determines a first phonetic representation of the at least one word (504), the first phonetic representation comprising a first set of phonemes selected from a speech recognition phonetic alphabet. In some implementations, the first phonetic representation of the at least one word is determined by a speech-to-text processor (e.g., the STT processing module 330, Figure 3A). The speech recognition phonetic alphabet is an alphabet or other symbolic linguistic representation used by the speech-to-text processor to recognize speech inputs. For example, in some implementations, the speech recognition phonetic alphabet is the International Phonetic Alphabet.

[0124] In some implementations, the speech-to-text processor determines the first phonetic representation by processing the speech input using an acoustic model to determine the phonemes in the utterance. For example, the set of phonemes for the word "Philippe" in above example above may be represented as "f-ill-ee-p-ay."

[0125] The digital assistant maps the first set of phonemes to a second set of phonemes to generate a second phonetic representation (e.g., with the phonetic alphabet conversion module 331, Figure 3A), wherein the second set of phonemes are selected from a speech

synthesis phonetic alphabet (506). In some implementations, the speech recognition phonetic alphabet is different than the speech synthesis phonetic alphabet. For example, the first set of phonemes in the speech recognition alphabet "f-ill-ee-p-ay" may be converted to the second set of phonemes "fill-ee-p-ay." By converting the first set of phonemes in a speech recognition alphabet to a speech synthesis alphabet, the pronunciation of the one or more words that were actually used by the user is capable of being used by both the speech recognizer and the speech synthesizer of the digital assistant, i.e., to recognize spoken words and to synthesize speech output.

[0126] In some implementations, the digital assistant includes conversion tables, maps, or equivalency lists to convert phonetic representations between different phonetic alphabets. Exemplary systems and methods for converting words and/or phonetic representations from one alphabet to another are described in Applicant's U.S. Utility Application Serial No. 13/469,047 for "Generalized Phonetic Transliteration Engine," filed May 10, 2012, the entire disclosure of which is incorporated herein by reference. In some implementations, the mapping includes using a weighted finite state transducer to map the first set of phonemes to the second set of phonemes. In some implementations, the mapping includes using a statistical model to map the first set of phonemes to the second set of phonemes.

[0127] The digital assistant stores the second phonetic representation in association with a text string corresponding to the at least one word (508). For example, the second phonetic representation "fill-ee-p-ay" is stored in a contact list in association with a contact having the name "Philippe" (e.g., as a user-specified pronunciation for a name in a contact list of a user). In some implementations, the second phonetic representation is stored in the user data 348 and/or the vocabulary index 344.

[0128] In some implementations, prior to receiving the first speech input, the text string is provided (501). In some implementations, the text string is a name in a contact list associated with a user. In some implementations, the name in the contact list is selected by a user, such as by navigating to a contact list application and selecting a name in the contact list for which to supply a preferred pronunciation. In some implementations, the text string is input by a user via a keyboard. For example, the user may type one or more words into a device (e.g., the user device 104) and then provide the first speech input (step 502) to specify a preferred pronunciation for one or more of the typed words. As another example, the user can select a contact (e.g., by pressing and holding a contact item in a contact list) and concurrently speak the preferred pronunciation for the name of the contact.

[0129] In some implementations, the text string is from a webpage displayed by an electronic device. For example, a user can select a word in a webpage (e.g., by touching, pressing, or clicking on the word), and then provide a speech input such as "pronounce this word as [word]."

[0130] Rather than requiring the user to manually identify the text string, the digital assistant may identify the text string automatically. In some implementations, the digital assistant determines the text string using the first phonetic representation (505). This may be accomplished by determining that the utterance corresponds to a certain sequence of letters, even if the digital assistant does not recognize that sequence of letters as a word. For example, a speech recognizer can determine that the phonemes "tuh-may-doe" correspond to the letters "t o m a t o," even if that word is not in the speech recognizer's vocabulary. In some implementations, the digital assistant uses fuzzy matching and/or approximate matching techniques to determine the text string from the first phonetic representation. For example, if a user provides a speech input to a digital assistant asking to call "f-ill-ee-p-ay," but this particular phonetic sequence has not been associated with the name "Philippe," the digital assistant uses fuzzy matching to determine that "f-ill-ee-p-ay" is sufficiently close to the sequence "f-ill-ee-p," which is a candidate pronunciation for the name Philippe. Thus, the digital assistant determines that the text string "Philippe" corresponds to the word in the utterance, even though the user's pronunciation of that word was not one of the candidate pronunciations.

[0131] In some implementations, the digital assistant updates a speech recognizer (e.g., the STT processing module 330) to associate the first phonetic representation with the text string (510). In some implementations, the first phonetic representation is stored as a candidate phonetic representation for the text string, for example, in a user-specific speech recognition vocabulary (e.g., the vocabulary index 344), or in a contact item in a contact list or address book (e.g., in the user data 348).

[0132] By associating the first phonetic representation of the word with the text string, the speech recognizer is able to identify the word when that particular pronunciation is used. For example, in some implementations, after updating the speech recognizer, the digital assistant receives a second speech input including the at least one word (512). In some implementations, the second speech input is any input received after the first input that includes the word. The digital assistant determines a third phonetic representation of the at least one word (514) (e.g., using the STT processing module 330), and determines that the at

Least one word corresponds to the text string based on a determination that the third phonetic representation is substantially similar to the first phonetic representation (516). Thus, the first phonetic representation reflecting the user's preferred pronunciation of a word is used to detect the word in later utterances.

[0133] Turning to Figure 5B, by storing the second phonetic representation (for speech synthesis) in association with the text string, the digital assistant is able to use the user-specified pronunciation in speech outputs that include the word. For example, in some implementations, after storing the second phonetic representation in association with the text string, the digital assistant synthesizes a speech output corresponding to the text string using the second phonetic representation (518). Accordingly, the synthesized speech output will sound substantially similar to the word in the speech input (e.g., the word as spoken by the user). As a specific example, after storing a second phonetic representation "fill-eep-ay" (corresponding to the user-specified pronunciation of the word in a speech synthesis phonetic alphabet), the digital assistant synthesizes a speech output using the user-specified pronunciation of the word "Philippe" (e.g., "Okay, I'm placing a telephone call to fill-eep-ay.")

[0134] In some implementations, the digital assistant stores a plurality of additional first phonetic representations in association with the text string (520), wherein the plurality of additional first phonetic representations includes the first phonetic representation. For example, the plurality of additional first phonetic representations may be the candidate phonetic representations described above. In particular, the first phonetic representation and a number of other candidate phonetic representations may be stored in association with the text string for use by a speech recognizer (e.g., the STT processing module 330) when identifying words in speech inputs. In some implementations, the candidate representations are specific to a particular user (e.g., they are stored in association with a particular user and/or user identifier), while in some implementations, they are not specific to any particular user (e.g., they are part of a generic vocabulary for use by multiple digital assistants and/or STT processors). By storing a plurality of first representations, the digital assistant provides more robust recognition functionality, because it will successfully recognizer more pronunciations of the word.

[0135] In some implementations, the digital assistant selects a predicted phonetic representation for the text string from the plurality of additional first phonetic representations (522). A predicted phonetic representation is one that is selected by the digital assistant as

the most probable pronunciation for that word. In some implementations, the predicted phonetic representation is a user-specified pronunciation. In some implementations, the predicted phonetic representation is a generic or global pronunciation for that word. In some implementations, all of the phonetic representations of the plurality of additional first phonetic representations are ranked in order of their likelihood of use, and the predicted phonetic representation ranked the highest.

[0136] In some implementations, where the candidate phonetic representations (including one or more user-specified pronunciations) are included in or are part of a generic vocabulary used by a plurality of digital assistants, the predicted phonetic representation is selected in for each user in accordance with a selection metric. In some implementations, the predicted phonetic representation is a most frequently detected phonetic representation.

[0137] In some implementations, the predicted phonetic representation is the most frequently detected phonetic representation for a particular region. For example, many users of digital assistants in a given region may specify the same pronunciation for the same word. Thus, the digital assistant (e.g., a digital assistant server or other centralized system that communicates with or otherwise interacts with many digital assistant clients) can determine the pronunciation that is most frequently specified by users in various geographical areas. Once a predicted phonetic representation is selected, it can be used for both speech recognition as well as speech synthesis (e.g., so that the digital assistant will recognize that pronunciation in speech inputs, as well as produce that pronunciation in speech outputs). The geographical region may be any appropriate geographical region(s), defined by any appropriate geographical, geopolitical, or cultural boundaries. For example, the predicted pronunciation for the word "tomato" may be "tuh-may-doe" for users in the Western portion of the United States, and "tuh-mah-doe" for users in the Eastern portion of the United States.

[0138] While the above discussion relates to selecting a predicted phonetic representation from a plurality of first representations (e.g., phonetic representations in a speech recognition phonetic alphabet), the discussion applies equally to selecting from among phonetic representations in any phonetic alphabet (e.g., phonetic representations in a speech synthesis phonetic alphabet).

[0139] Figure 6 is a flow diagram of an exemplary method 600 implemented by a digital assistant for learning word pronunciations. In some implementations, the method 600 is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. For example, in some

implementations, the method 600 is performed at the user device 104 or the server system 108. In some implementations, the method 600 is performed by the digital assistant system 300 (Figure 3A), which, as noted above, may be implemented on a standalone computer system (e.g., either the user device 104 or the server system 108) or distributed across multiple computers (e.g., the user device 104, the server system 108, and/or additional or alternative devices or systems). While the following discussion describes the method 600 as being performed by a digital assistant (e.g., the digital assistant system 300), the method is not limited to performance by any particular device or combination of devices. Moreover, the individual steps of the method may be distributed among the one or more computers, systems, or devices in any appropriate manner.

[0140] The method 600 relates to detecting errors in speech based interactions in order to determine that the digital assistant should request user-specified pronunciations for one or more words. Various types of errors may arise in speech based interactions. For example, errors in speech recognition can prevent the digital assistant from accurately inferring the user's intent. Specifically, a STT processor might select the wrong word based on a speech input (e.g., detecting the word "potato" when the user actually said "tomato"), or might simply fail to identify a suitable word for a portion of a text string (e.g., failing to transcribe a last name of a contact). Moreover, the digital assistant might mispronounce a word (such as a contact name) during a speech interaction, which can confuse the user. This is especially problematic when the digital assistant is used in an "eyes-free" mode, where the user cannot simply read a text version of the speech output on a display screen to determine whether the digital assistant has identified the wrong name, or simply mispronounced a correctly identified name. For example, while driving, a user may say to a digital assistant "Call Philippe," and the digital assistant responds by saying "Calling [fill-up]." Upon hearing the mispronounced name, the user may not trust that the digital assistant is placing a call to the correct person, and may unnecessarily cancel the telephone call. Accordingly, when an error is detected, it is often possible for the digital assistant to infer that the error was likely caused by either an inability to recognize a user's pronunciation of a word, or an incorrect pronunciation in a speech output by the digital assistant. In some implementations, errors are detected automatically by the digital assistant (e.g., based on certain patterns of use or interactions with the digital assistant or a user device that are indicative of an error), or are manually indicated by a user (e.g., when a user provides a speech input that indicates an

error, or selects an affordance indicating that an error has occurred). Various specific ways of detecting errors are described herein.

[0141] Returning to the method 600, the digital assistant detects an error in a speech based interaction with a digital assistant (608). In some implementations, the error is an error in speech recognition of one or more words that were received in a speech input provided by the user. In some implementations, the error is an error in speech synthesis of one or more words that were output in a speech output by the electronic device (e.g., a mispronunciation of one or more words). Specific examples of errors in speech recognition and synthesis, and how they are detected, are provided below.

[0142] In response to detecting the error, the digital assistant receives a speech input from a user, the speech input including a pronunciation of one or more words (612). In some implementations, prior to receiving the speech input from the user, the digital assistant prompts the user to provide a speech input including a preferred pronunciation after detecting an error (610). For example, in some implementations, after detecting an error in recognizing the name "Philippe," the digital assistant will display or say "I'm sorry I didn't understand you just now - can you please tell me how you pronounce "Philippe" that so I can remember it?" The user can then speak the word to the digital assistant.

[0143] In some implementations, when the digital assistant prompts the user to provide a pronunciation of one or more words, the digital assistant only requests a pronunciation of one word at a time, so as to avoid detecting or incorporating erroneous phonemes from other words in the utterance. Specifically, if the user provides a speech input such as "please pronounce that as 'Philippe' from now on," the speech recognizer may have difficulty parsing what syllables should be associated with the name "Philippe."

[0144] In some implementations, when user-specified pronunciations are requested for more than one word (e.g., a first name and a last name of a contact item), the digital assistant provides separate prompts, and receives separate inputs, for each name. Specifically, the digital assistant may say "How shall I pronounce 'Philippe'?" After receiving a response, the digital assistant may then say "Thanks. How shall I pronounce 'Martin'?"

[0145] In some implementations, the digital assistant determines a first phonetic representation of the one or more words (614), the first phonetic representation comprising a first set of phonemes selected from a speech recognition phonetic alphabet. In some implementations, the first phonetic representation of the at least one word is determined by a speech-to-text processor (e.g., the STT processing module 330, Figure 3A).

[0146] In some implementations, the digital assistant maps the first set of phonemes to a second set of phonemes to generate a second phonetic representation (e.g., with the phonetic alphabet conversion module 331, Figure 3A), wherein the second set of phonemes are selected from a speech synthesis phonetic alphabet (616). In some implementations, the speech recognition phonetic alphabet is different than the speech synthesis phonetic alphabet. Mapping phonemes from one phonetic alphabet to another is discussed in greater detail with respect to step (506) in Figure 5A, above.

[0147] The digital assistant stores the pronunciation (including either or both of the first phonetic representation and the second phonetic representation) in association with a text string corresponding to the one or more words (614). In some implementations, receiving and storing the pronunciation in association with the text string includes performing one or more of the steps of the method 500, described above.

[0148] One type of error that the digital assistant detects is a low confidence in the recognition of a word in a speech input. Thus, in some implementations, the digital assistant receives (prior to detecting the error at step (608)) a speech input including one or more words (602). The digital assistant performs speech recognition on the speech input (e.g., with the STT processing module 330) to generate a text string corresponding to the one or more words (604). The digital assistant determines a confidence metric of the text string (606). In some implementations, the confidence metric reflects how closely the phoneme set generated for a particular word (e.g., by the STT processing module 330) corresponds to one or more of the candidate phoneme sets for that word. The digital assistant then detects the error (608) based on a determination that the confidence metric does not meet a predetermined threshold.

[0149] The digital assistant can also detect errors in speech synthesis, such as pronunciation errors made by the digital assistant in a synthesized speech output. In some implementations, the digital assistant synthesizes (prior to detecting the error at step (608)) a speech output including one or more words (607). For example, the speech output may be used to indicate to the user what action the digital assistant is taking in response to a request from the user, such as when the digital assistant says "Calling Philippe Martin" in response to a request to initiate a phone call to that person. The digital assistant then detects the error (608) based on an indication from the user that the one or more words were pronounced incorrectly by the digital assistant. In some implementations, the indication from the user is a speech input where the user indicates that the digital assistant's pronunciation was incorrect, such as "You said that name wrong" or "That's pronounced [f-ill-ee-p-ay]." In some

implementations, the indication is a selection of an affordance to cancel a telephone call that was initiated by the digital assistant.

[0150] In some implementations, the digital assistant detects the error (608) based on detecting a user selection of the affordance. An affordance that a user can select to indicate an error may be displayed to the user during interactions with the digital assistant. For example, in some implementations, when providing a synthesized speech output, the digital assistant also displays a textual version of the speech output to the user on a display. This may be helpful, for example, so that if a word is mispronounced by the digital assistant, the user can simply refer to the displayed text of the output to determine what the digital assistant was trying to speak. In some implementations, the digital assistant provides the affordance (e.g., a touchscreen button, a physical button, etc.) along with the displayed text so that the user can indicate that a speech output includes a mispronounced word. In some implementations, the affordance is labeled "Did I mispronounce something?" or "Click here to correct pronunciation mistakes."

[0151] In some implementations, the digital assistant also determines confidence values for speech synthesis, and determines whether a speech output is likely to be misunderstood by a user and/or be likely to cause a user to correct the digital assistant's pronunciation. In particular, speech synthesizers may be very confident that it will correctly pronounce common words and/or words that are in a vocabulary of known (e.g., pre-recorded) pronunciations. For example, speech synthesizers will be very confident that it will correctly pronounce the words "call" and "the" and "sure," among a multitude of possible examples. However, for words that are not in the vocabulary or for which no user-specified pronunciation has been provided, the digital assistant can identify that its pronunciation of those words are more likely to be incorrect. Accordingly, in some implementations, the digital assistant determines confidence values for words in speech synthesis outputs. The digital assistant can then monitor user interactions with the digital assistant and/or a user device to infer whether the pronunciation was correct or not. For example, if the digital assistant receives an input such as "what was that?" after synthesizing one or more words with a low pronunciation confidence value, the digital assistant responds (either immediately or at a later time) with a response that better informs the user of the intended word (either by spelling it or displaying it), requests correction by the user, or both. For example, the digital assistant may provide an output such as "Sorry, I may have said that wrong. Here's what I was trying to say p-h-i-l-i-p-p-e. Can you tell me how to pronounce that?" Providing a

speech output that spells the word having a low confidence score is also beneficial when the digital assistant is being used in an eyes-free mode, where the user cannot (or should not) look at a display to read text. In particular, the user can easily and quickly understand what the digital assistant was trying to say, and the dialogue between the user and the digital assistant can continue once the pronunciation error is resolved (or dismissed by the user).

[0152] In some implementations, when an error is detected in a speech output (e.g., by detecting a user interaction indicative of a mistake after synthesizing a word with a low pronunciation confidence level, as described above), the digital assistant does not immediately prompt the user to provide a correct pronunciation, but instead spells the word for the user and attempts to continue to take appropriate actions to satisfy the user's intent. For example, if, after synthesizing the phrase "Calling Philippe Martin," the digital assistant receives an input such as "who?", the digital assistant may respond by saying "Sorry, I must have mispronounced that. Would you like me to call [philippe] Martin?" In some implementations, after the user's intent is satisfied (or if the user dismisses or cancels the action), the digital assistant will request that the user provide a user-specified pronunciation for the word that it believes to have been mispronounced.

[0153] Figure 7 is a flow diagram of an exemplary method 700 implemented by a digital assistant for learning word pronunciations. In some implementations, the method 700 is performed at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors. For example, in some implementations, the method 700 is performed at the user device 104 (or a plurality of user devices) or the server system 108. In some implementations, the method 700 is performed by the digital assistant system 300 (Figure 3A), which, as noted above, may be implemented on a standalone computer system (e.g., either the user device 104 or the server system 108) or distributed across multiple computers (e.g., the user device 104, the server system 108, and/or additional or alternative devices or systems). While the following discussion describes the method 700 as being performed by a digital assistant (e.g., the digital assistant system 300), the method is not limited to performance by any particular device or combination of devices. Moreover, the individual steps of the method may be distributed among the one or more computers, systems, or devices in any appropriate manner.

[0154] In accordance with some implementations, the method 700 allows a pronunciation specified by one user to be used by other users for both speech recognition and synthesis. For example, once a user of a digital assistant and/or a digital assistant service specifies how his

or her own name is pronounced, that pronunciation can be provided to other users' digital assistants and/or devices (with appropriate permissions by the first user) so that those digital assistants will recognize and pronounce the name correctly. In some implementations, a user's own name pronunciation is only shared if the user specifically authorizes it.

[0155] In some implementations, a user-specified pronunciation for a first user's name is shared with a second user (subject to authorization) when the second user creates a contact item (e.g., in a contact list or address book) for the first user. The contact item may be provided to the second user by the first user directly (e.g., by sharing an electronic business card or contact file), or may be manually created by the second user (e.g., when the second user types in a name, phone number, or email address, etc., of the user into a contact list or address book on the second user's device).

[0156] Returning to Figure 7 and method 700, the digital assistant receives from a first user, a user-specified pronunciation of the first user's name (702). The digital assistant stores the pronunciation of the first user's name in association with the first user's name and a unique identifier of the first user (704). In some implementations, the first user's name is a textual representation of the first user's name. In some implementations, the unique identifier is an email address of the first user. In some implementations, the unique identifier is a telephone number of the first user.

[0157] In some implementations, the first user's name and the unique identifier of the first user are stored in an electronic business card. In some implementations, the electronic business card is in a vCard format.

[0158] In some implementations, the pronunciation, the name (e.g., a textual representation of the name), and the unique identifier are stored on a device associated with the first user (e.g., the user device 104). In some implementations, the pronunciation, the name, and the unique identifier are stored on a central server system associated with digital assistants of a plurality of users (e.g., the server system 108). In some implementations, the pronunciation, the name, and the unique identifier are stored on a central server system and on the user devices of one or more individual users (including, for example, the first user and other users authorized by the first user to access the pronunciation, as discussed below).

[0159] The digital assistant detects that a second user has created a contact item including the first user's name and the unique identifier of the first user (706). The digital assistant uses the user-specified pronunciation for one or both of recognizing the first user's name in

speech inputs by the second user, and synthesizing the first user's name in speech outputs to the second user by the digital assistant (708).

[0160] In some implementations, when the second user creates and/or stores a new contact item, the second user's device communicates with the server system to determine if the user identified by the contact item is associated with a user-specified pronunciation, and, if so, whether the second user is authorized to access it. In some implementations, if the second user is not authorized to access the user-specified pronunciation, an authorization request is provided to the first user to authorize the second user to receive and/or access the user-specified pronunciation. In some implementations, the second user is not made aware that it is accessing a user-specified pronunciation of the first user's name. Thus, once the second user's digital assistant gains access to the user-specified pronunciation, it can be adopted and used by the second user's digital assistant, making it more accurate and making it appear even more intelligent and useful.

[0161] While the method 700 is described above for user-specified pronunciations of names, the method also applies to other words as well. For example, a first user may specify a pronunciation of a business name, city, street name, or the like. These user-specified pronunciations can then be used by other users' digital assistants (subject to authorization) to correctly recognize and pronounce these words in speech inputs and outputs.

[0162] The operations described above with reference to Figures 5A-7 are, optionally, implemented by components depicted in Figure 2 and/or Figure 3. Similarly, it would be clear to a person having ordinary skill in the art how other processes can be implemented based on the components depicted in Figure 2 and/or Figure 3.

[0163] It should be understood that the particular order in which the operations have been described above is merely exemplary and is not intended to indicate that the described order is the only order in which the operations could be performed. One of ordinary skill in the art would recognize various ways to reorder the operations described herein.

[0164]

[0165] The foregoing description, for purpose of explanation, has been described with reference to specific implementations. However, the illustrative discussions above are not intended to be exhaustive or to limit the invention to the precise forms disclosed. Many modifications and variations are possible in view of the above teachings. The implementations were chosen and described in order to best explain the principles of the invention and its practical applications, to thereby enable others skilled in the art to best utilize the invention and various implementations with various modifications as are suited to the particular use contemplated.

What is claimed is:

1. A method for learning word pronunciations, comprising:
 - at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors:
 - receiving a first speech input including at least one word;
 - determining a first phonetic representation of the at least one word, the first phonetic representation comprising a first set of phonemes selected from a speech recognition phonetic alphabet;
 - mapping the first set of phonemes to a second set of phonemes to generate a second phonetic representation, the second set of phonemes selected from a speech synthesis phonetic alphabet; and
 - storing the second phonetic representation in association with a text string corresponding to the at least one word.
2. The method of claim 1, further comprising, prior to receiving the first speech input, providing the text string.
3. The method of claim 2, wherein the text string is a name in a contact list associated with a user.
4. The method of claim 2, wherein the text string is input by a user via a keyboard.
5. The method of claim 2, wherein the text string is from a webpage displayed by the electronic device.
6. The method of claim 1, further comprising determining the text string using the first phonetic representation.
7. The method of claim 1, further comprising updating a speech recognizer to associate the first phonetic representation with the text string.
8. The method of claim 7, further comprising:
 - after updating the speech recognizer, receiving a second speech input including the at least one word;
 - determining a third phonetic representation of the at least one word; and

determining that the at least one word corresponds to the text string based on a determination that the third phonetic representation is substantially similar to the first phonetic representation.

9. The method of claim 1, further comprising, after storing the second phonetic representation in association with the text string, synthesizing a speech output corresponding to the text string using the second phonetic representation.

10. The method of claim 9, wherein the synthesized speech output of the text string sounds substantially similar to the at least one word of the first speech input.

11. The method of claim 7, further comprising:

storing a plurality of additional first phonetic representations in association with the text string, wherein the plurality of additional first phonetic representations includes the first phonetic representation; and

selecting a predicted phonetic representation for the text string from the plurality of additional first phonetic representations.

12. The method of claim 11, wherein the predicted phonetic representation is a most frequently detected phonetic representation in the plurality of additional first phonetic representations.

13. The method of claim 11, wherein the predicted phonetic representation is the most frequently detected phonetic representation for a particular geographical region.

14. The method of claim 1, wherein the mapping includes using a weighted finite state transducer to map the first set of phonemes to the second set of phonemes.

15. The method of claim 1, wherein the mapping includes using a statistical model to map the first set of phonemes to the second set of phonemes.

16. The method of claim 1, wherein the speech recognition phonetic alphabet is different than the speech synthesis phonetic alphabet.

17. A method for learning word pronunciations, comprising:

at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors:

detecting an error in a speech based interaction with a digital assistant;
in response to detecting the error, receiving a speech input from a user, the speech input including a pronunciation of one or more words; and
storing the pronunciation in association with a text string corresponding to the one or more words.

18. The method of claim 17, wherein the one or more words were received in a prior speech input provided by the user, and wherein the error is an error in speech recognition of the one or more words.

19. The method of claim 17, wherein the one or more words were output in a speech output by the electronic device, and wherein the error is an error in speech synthesis of the one or more words.

20. The method of claim 17, further comprising:
receiving a speech input including the one or more words;
performing speech recognition on the speech input to generate a text string corresponding to the one or more words;
determining a confidence metric of the text string; and
detecting the error based on a determination that the confidence metric does not meet a predetermined threshold.

21. The method of claim 17, further comprising:
synthesizing a speech output including the one or more words; and
detecting the error based on an indication from the user that the one or more words were pronounced incorrectly.

22. The method of claim 17, further comprising, after detecting the error, prompting the user to provide the speech input.

23. A method for learning word pronunciations, comprising:
at an electronic device with one or more processors and memory storing one or more programs for execution by the one or more processors:
receiving, from a first user, a user-specified pronunciation of the first user's name;

storing the pronunciation of the first user's name in association with the first user's name and a unique identifier of the first user;

detecting that a second user has created a contact item including the first user's name and the unique identifier of the first user; and


using the user-specified pronunciation for one or both of recognizing the first user's name in speech inputs by the second user, and synthesizing the first user's name in speech outputs to the second user by the digital assistant.

24. The method of claim 23, wherein the unique identifier is an email address of the first user.
25. The method of claim 23, wherein the unique identifier is a telephone number of the first user.
26. The method of claim 23, wherein the first user's name and the unique identifier of the first user are stored in an electronic business card.
27. The method of claim 26, wherein the electronic business card is in a vCard format.
28. A method comprising features of any combination of the methods of claims 1-27.
29. A system comprising one or more electronic devices, each comprising one or more of:
 - a display;
 - a touch-sensitive surface;
 - one or more processors;
 - memory; andone or more programs, wherein 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 performing any of the methods of claims 1-27.
30. A computer readable storage medium storing one or more programs, the one or more programs comprising instructions, which when executed by an electronic device with a display, a touch-sensitive surface cause the device to perform any of the methods of claims 1-27.
31. An electronic device, comprising:

a display;
a touch-sensitive surface; and
means for performing any of the methods of claims 1-27.

32. An information processing apparatus for use in an electronic device with a display, a touch-sensitive surface, comprising:
means for performing any of the methods of claims 1-27.

FIG. 1






FIG. 2





FIG. 3A

FIG. 3B






FIG. 3C




FIG. 4




FIG. 5A




FIG. 5B




FIG. 6

10/10




FIG. 7