



(51) International Patent Classification:
G10L 15/00 (2006.01)

(21) International Application Number:
PCT/US2010/024895

(22) International Filing Date:
22 February 2010 (22.02.2010)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
12/389,678 20 February 2009 (20.02.2009) US

(71) Applicant (for all designated States except US): **VOICE-BOX TECHNOLOGIES, INC.** [US/US]; 11980 NE 24th Street, Suite 100, Bellevue, WA 98005 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): **BALDWIN, Larry** [US/US]; 23125 SE 243rd Place, Maple Valley, WA 98038 (US). **WEIDER, Chris** [US/US]; 10757 Greenwood Avenue North #220, Seattle, WA 98133 (US).

(74) Agents: **ALI, Syed, Jafar** et al.; Pillsbury Winthrop Shaw Pittman LLP, P.O. Box 10500, McLean, VA 22102 (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, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

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

Published:

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

(54) Title: SYSTEM AND METHOD FOR PROCESSING MULTI-MODAL DEVICE INTERACTIONS IN A NATURAL LANGUAGE VOICE SERVICES ENVIRONMENT

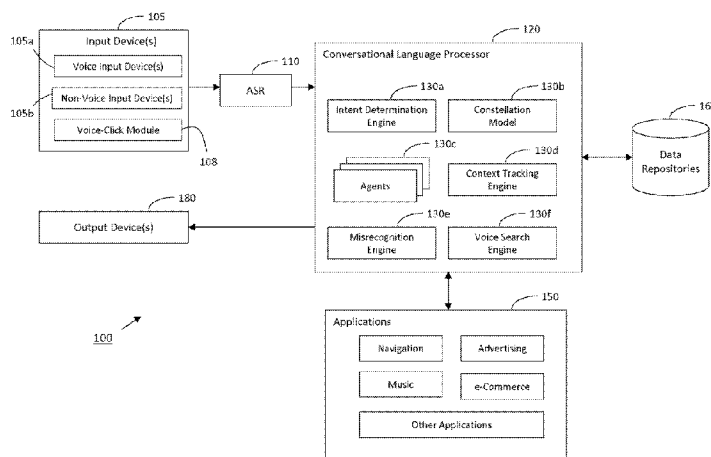


Figure 1

(57) Abstract: A system and method for processing multi-modal device interactions in a natural language voice services environment may be provided. In particular, one or more multi-modal device interactions may be received in a natural language voice services environment that includes one or more electronic devices. The multi-modal device interactions may include a non-voice interaction with at least one of the electronic devices or an application associated therewith, and may further include a natural language utterance relating to the non-voice interaction. Context relating to the non-voice interaction and the natural language utterance may be extracted and combined to determine an intent of the multi-modal device interaction, and a request may then be routed to one or more of the electronic devices based on the determined intent of the multi-modal device interaction.

SYSTEM AND METHOD FOR PROCESSING MULTI-MODAL DEVICE INTERACTIONS IN A NATURAL LANGUAGE VOICE SERVICES ENVIRONMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

[001] This application claims the benefit of U.S. Patent Application Serial No. 12/389,678, entitled "SYSTEM AND METHOD FOR PROCESSING MULTI-MODAL DEVICE INTERACTIONS IN A NATURAL LANGUAGE VOICE SERVICES ENVIRONMENT", filed February 20, 2009, which is hereby incorporated by reference in its entirety.

FIELD OF THE INVENTION

[002] The invention relates to an integrated natural language voice services environment for processing multi-modal interactions with or one or more devices and/or applications, wherein the multi-modal interactions may provide additional context for cooperatively interpreting and otherwise processing a natural language utterance accompanying the multi-modal interactions.

BACKGROUND OF THE INVENTION

[003] As technology has progressed in recent years, consumer electronic devices have emerged to become nearly ubiquitous in the everyday lives of many people. To meet the increasing demand that has resulted from growth in the functionality and mobility of mobile phones, navigation devices, embedded devices, and other such devices, many devices offer a wealth of features and functions in addition to core applications. Greater functionality also introduces trade-offs, however, including learning curves that often inhibit users from fully exploiting all of the capabilities of their electronic devices. For example, many existing electronic devices include complex human to machine interfaces that may not be particularly user-friendly, which can inhibit mass-market adoption for many technologies. Moreover, cumbersome interfaces often result in otherwise desirable features being difficult to find or use (e.g., because of menus that are complex or otherwise tedious to navigate). As such, many users tend not to use, or even know about, many of the potential capabilities of their devices.

[004] As such, the increased functionality of electronic devices often tends to be wasted, as market research suggests that many users only use only a fraction of the features or applications available on a given device. Moreover, in a society where wireless networking and broadband access are increasingly prevalent, consumers tend to naturally desire seamless

mobile capabilities from their electronic devices. Thus, as consumer demand intensifies for simpler mechanisms to interact with electronic devices, cumbersome interfaces that prevent quick and focused interaction become an important concern. Nevertheless, the ever-growing demand for mechanisms to use technology in intuitive ways remains largely unfulfilled.

[005] One approach towards simplifying human to machine interactions in electronic devices has included the use of voice recognition software, which has the potential to enable users to exploit features that would otherwise be unfamiliar, unknown, or difficult to use. For example, a recent survey conducted by the Navteq Corporation, which provides data used in a variety of applications such as automotive navigation and web-based applications, demonstrates that voice recognition often ranks among the features most desired by consumers of electronic devices. Even so, existing voice user interfaces, when they actually work, still require significant learning on the part of the user.

[006] For example, many existing voice user interface only support requests formulated according to specific command-and-control sequences or syntaxes. Furthermore, many existing voice user interfaces cause user frustration or dissatisfaction because of inaccurate speech recognition. Similarly, by forcing a user to provide pre-established commands or keywords to communicate requests in ways that a system can understand, existing voice user interfaces do not effectively engage the user in a productive, cooperative dialogue to resolve requests and advance a conversation towards a satisfactory goal (e.g., when users may be uncertain of particular needs, available information, device capabilities, etc.). As such, existing voice user interfaces tend to suffer from various drawbacks, including significant limitations on engaging users in a dialogue in a cooperative and conversational manner.

[007] Additionally, many existing voice user interfaces fall short in utilizing information distributed across different domains, devices, and applications in order to resolve natural language voice-based inputs. Thus, existing voice user interfaces suffer from being constrained to a finite set of applications for which they have been designed, or to devices on which they reside. Although technological advancement has resulted in users often having several devices to suit their various needs, existing voice user interfaces do not adequately free users from device constraints. For example, users may be interested in services associated with different applications and devices, but existing voice user interfaces tend to restrict users from accessing the applications and devices as they see fit. Moreover, users typically can only practicably carry

a finite number of devices at any given time, yet content or services associated with users' other devices currently being used may be desired in various circumstances.

[008] Accordingly, although users tend to have varying needs, where content or services associated with different devices may be desired in various contexts or environments, existing voice technologies tend to fall short in providing an integrated environment in which users can request content or services associated with virtually any device or network. As such, constraints on information availability and device interaction mechanisms in existing voice services environments tend to prevent users from experiencing technology in an intuitive, natural, and efficient way. For instance, when a user wishes to perform a given function using a given electronic device but does not necessarily know how to go about performing the function, the user typically cannot engage in a multi-modal interaction with the device to simply utter words in natural language to request the function.

[009] Furthermore, relatively simple functions can often be tedious to perform using electronic devices that do not have voice recognition capabilities. For example, purchasing new ring-tones for a mobile phone tends to be a relatively straightforward process, but users must typically navigate several menus and press many different buttons in order to complete the process. As such, it becomes apparent that interaction with electronic devices could be far more efficient if users were able to use natural language to exploit buried or otherwise difficult to use functionality. Existing systems suffer from these and other problems.

SUMMARY OF THE INVENTION

[010] According to one aspect of the invention, a system and method for processing multi-modal device interactions in a natural language voice services environment may be provided. In particular, one or more multi-modal interactions may be received in a natural language voice services environment that includes one or more electronic devices. The multi-modal device interactions may include a user engaging in a non-voice interaction with one or more of the electronic devices or applications associated with the devices, while also providing a natural language utterance in relation to the non-voice interaction. For example, the non-voice device interaction may comprise the user selecting a particular segment, item, data, point of focus, or attention focus, or otherwise engaging in one or more unique and distinguishable interactions with the electronic devices or applications associated therewith. As such, context may be

extracted from the natural language utterance, and the non-voice device interaction may provide additional context for the natural language utterance. The context of the utterance and the non-voice device interaction may then be combined to determine an intent of the multi-modal device interaction, wherein one or more of the electronic devices may process a request based on the intent of the multi-modal device interaction.

[011] According to one aspect of the invention, at least one of the electronic devices may include an input device configured to receive voice-based inputs. In one implementation, the voice-based input device may be signaled to capture the natural language utterance in response to detecting the non-voice interaction with the one or more electronic devices or applications. Furthermore, the natural language voice services environment may include one or more listeners established for the electronic devices and the associated applications, wherein the listeners may be configured to detect the non-voice interaction with the electronic devices or applications. As such, information relating to the non-voice interaction and the accompanying natural language utterance may be aligned to enable cooperative processing of the utterance and the non-voice device interaction.

[012] According to one aspect of the invention, at least one transaction lead may be generated based on the intent of the multi-modal device interaction. For example, a further multi-modal device interaction may be received, wherein the further multi-modal device interaction may relate to the transaction lead generated for the first multi-modal device interaction. At least one request may then be routed to one or more of the electronic devices based on an intent determined for the further multi-modal device interaction, whereby a transaction click-through may be processed in response to receiving the device interaction relating to the generated transaction lead. For example, the transaction lead may include an advertisement or a recommendation that is selected based on the intent of the original multi-modal device interaction, while the further multi-modal device interaction may include a user selecting the advertisement or recommendation. Thus, the selection of the advertisement or recommendation may be considered a transaction click-through, which may generate revenue for a particular entity (e.g., a provider of the natural language voice services environment).

[013] Other objects and advantages of the invention will be apparent based on the following drawings and detailed description.

BRIEF DESCRIPTION OF THE DRAWINGS

[014] Figure 1 illustrates a block diagram of an exemplary system for processing multi-modal device interactions in a natural language voice services environment, according to various aspects of the invention.

[015] Figure 2 illustrates a block diagram of an exemplary method for synchronizing multi-modal devices in a natural language voice service environment, according to various aspects of the invention.

[016] Figure 3 illustrates a flow diagram of an exemplary method for processing multi-modal device interactions in a natural language voice service environment, according to various aspects of the invention.

[017] Figure 4 illustrates a flow diagram of an exemplary method for processing multi-modal device interactions to generate one or more transaction leads in a natural language voice service environment, according to various aspects of the invention.

DETAILED DESCRIPTION

[018] According to various aspects of the invention, Figure 1 illustrates a block diagram of an exemplary system 100 for processing multi-modal device interactions in a natural language voice services environment. As will be apparent from the further description to be provided herein, the system 100 illustrated in Figure 1 may include an input device 105, or a combination of input devices 105, which enable users to interact with the system 100 in a multi-modal manner. In particular, the system 100 may include various natural language processing components, including at least a voice-click module 108, which can collectively process a user's multi-modal interactions with one or more input devices 105. For example, in one implementation, the input devices 105 may include any suitable combination of at least one voice input device 105a (e.g., a microphone) and at least one non-voice input device 105b (e.g., a mouse, touch-screen display, wheel selector, etc.). As such, the input devices 105 may include any suitable combination of electronic devices having mechanisms for receiving both voice-based and non-voice-based inputs (e.g., a microphone coupled to one or more of a telematics device, personal navigation device, mobile phone, VoIP node, personal computer, media device, embedded device, server, or other electronic device). As such, the system 100 may enable users to engage in multi-modal conversational interactions with the one or more

electronic input devices 105 or applications associated with the electronic devices 105, wherein the system 100 may process the device interactions in a free-form and cooperative manner suitable for routing tasks or otherwise resolving requests.

[019] As indicated above, in one implementation, the system may include various natural language processing components that can support free-form utterances and/or other forms of device interactions, which may liberate users from restrictions relating to the manner of formulating commands, queries, or other requests. As such, users may interact with the input devices 105 using any manner of speaking into the voice input device 105a or otherwise interacting with the non-voice input devices 105b in order to request content or services available in the system 100. For instance, a user may request any content or services that may be available in the system 100 by providing a natural language utterance to the voice input device 105a. In one implementation, the utterance may then be processed using techniques described in U.S. Patent No. 7,398,209, entitled "Systems and Methods for Responding to Natural Language Speech Utterance," issued July 8, 2008, and U.S. Patent Application 10/618,633, entitled "Mobile Systems and Methods for Responding to Natural Language Speech Utterance," filed June 15, 2003, the disclosures of which are hereby incorporated by reference in their entirety. In addition, the user may interact with one or more of the non-voice input devices 105b to provide further context or other information relating to the utterance and/or the requested content or services.

[020] In one implementation, the system 100 may be coupled to various other systems that include additional multi-modal devices, with the other systems having natural language processing capabilities similar to those shown in Figure 1. The system 100 may therefore provide an interface to a multi-device environment in which users may request content or services available through the various additional devices in the environment. For example, in one implementation, the system 100 may include a constellation model 130b that provides knowledge relating to content, services, applications, intent determination capabilities, and other features available through the other systems and devices in the environment. For example, in one implementation, the system 100 may interact with devices, applications, or other systems in the environment to cooperatively resolve requests, as described in co-pending U.S. Patent Application 12/127,343, entitled "System and Method for an Integrated, Multi-Modal, Multi-Device Natural Language Voice Services Environment," filed May 27, 2008, the

disclosure of which is hereby incorporated by reference in its entirety. For example, the multi-device environment may share information among various systems and devices to provide a cooperative environment for resolving requests, wherein the shared information may relate to aspects such as device capabilities, context, prior interactions, domain knowledge, short-term knowledge, long-term knowledge, and cognitive models, among other things.

[021] As indicated above, the system 100 illustrated in Figure 1 may include, among other things, one or more electronic input devices 105 that collectively provide an interface (or combination of interfaces) for receiving one or more multi-modal device interactions from a user, wherein the device interactions include at least a user-spoken utterance. Although the implementation illustrated in Figure 1 includes a distinct voice input device 105a and non-voice input device 105b, it will be apparent that in one or more implementations the voice input device 105a and the non-voice input device 105b may be components of the same or separate devices. For example, the input devices 105 may include a microphone coupled to a mobile phone (i.e., the voice input device 105a), and may further include one or more buttons, selectable displays, wheel selectors, or other components coupled to the mobile phone (i.e., the non-voice input devices 105b). In another example, the input devices 105 may include a combination of a microphone coupled to a telematics device (i.e., the voice input device 105a), and may further include buttons, a touch-screen display, a track wheel, or other non-voice input devices 105b coupled to a media player that is communicatively coupled to, yet distinct from, the telematics device. Thus, the input devices 105 may include any suitable combination of communicatively coupled electronic devices that includes at least one input device for receiving natural language utterance inputs and at least one input device for receiving multi-modal non-voice inputs.

[022] In one implementation, a voice-click module 108 communicatively coupled to the one or more input devices 105 may enable cooperative processing of multi-modal device interactions received at the voice input device 105a and one or more of the non-voice input devices 105b. For example, the voice-click module 108 may provide the system 100 with information that can be used to process a natural language utterance received via the voice input device 105a in view of one or more non-voice device interactions received via the non-voice input devices 105b. The voice-click module 108 therefore enables the user to interact with the various input devices 105 in an intuitive and free-form manner, whereby the user may provide various types

of information to the system 100 when seeking to initiate action, retrieve information, or otherwise request content or services available in the system 100.

[023] The voice input device 105a may comprise any appropriate device, or combination of devices, which have capabilities for receiving natural language utterances or other forms of spoken input. For example, in one implementation, the voice input device 105a may include a directional microphone, an array of microphones, or other devices capable of creating encoded speech. In one implementation, the voice input device 105a may be configured to maximize fidelity of the encoded speech. For example, the voice input device 105a may be configured to maximize gain in a direction of the user, cancel echoes, null point noise sources, perform variable rate sampling, filter environmental noise or background conversations, or use other techniques to maximize the fidelity of the encoded speech. As such, the voice input device 105a may create the encoded speech in a manner tolerant of noise or other factors that could otherwise interfere with the system 100 accurately interpreting natural language utterances.

[024] The non-voice input device 105b may comprise any appropriate device, or combination of devices, which have capabilities for supporting non-voice device interactions. For example, in one implementation, the non-voice input device 105b may include a combination stylus and touch-screen or tablet interface, a BlackBerry® wheel selector, an iPod® click wheel, a mouse, a keypad, buttons, or any other devices that supports distinguishable non-voice device interactions. The user may therefore use the non-voice input devices 105b to make data selections or identify a point of focus (or attention focus) to be processed in connection with a related natural language utterance provided via the voice input device 105a. For example, the user may point a stylus at a specific segment of a touch-screen display, highlight text using a mouse, click a button, interact with an application, or otherwise engage in any suitable device interaction for selecting data or otherwise identifying a point of focus (i.e., voice-activating or “voice-clicking” the selected data and/or the identified point of focus).

[025] Furthermore, in addition to being usable for making data selections, identifying points of focus, or otherwise activating data to be interpreted in association with one or more utterances, the user may further use the non-voice input devices 105b to engage in specialized device interactions that have meaning in the system 100. For example, the specialized device interactions (which may be referred to as “clicks” or “voice-clicks”) may include clicks lasting a given duration, clicks continuously held for a given duration, clicks made in a predetermined

sequence, or any other interaction or sequence of interactions that the input devices 105 and/or the voice-click module 108 can identify, detect, or otherwise distinguish.

[026] In one implementation, the specialized device interactions may be associated with one or more actions, queries, commands, tasks, or other requests associated with applications or services available in the system 100. In one implementation, the specialized device interactions may further include one or more actions, queries, commands, tasks, or other requests associated with any of various devices deployed in the multi-device environment, as described in the above-referenced co-pending U.S. Patent Application 12/127,343, entitled "System and Method for an Integrated, Multi-Modal, Multi-Device Natural Language Voice Services Environment," filed May 27, 2008. For example, a distinct sequence of clicking a stylus on a particular segment or item displayed on a touch-screen display may be defined as a specialized device interaction or voice-click for initiating a telephone call on a mobile phone, calculating a route on a navigation device, purchasing a song to a media player, or another type of request.

[027] Thus, the voice-click module 108 coupled to the input device 105 may continually monitor the user's interactions with the non-voice input device 105b to detect occurrences of at least one non-voice device interaction, which may be referred to herein as a "voice-click." The detected voice-click may therefore provide further context for processing a multi-modal device interaction, which may include the at least one voice-click and one or more natural language utterances, each of which may provide context for task specification. Thus, the voice-click may generally signal the system 100 that a current utterance or other voice-based input is to be processed together with a current interaction with one or more of the devices 105. For example, in one implementation, the current device interaction may include the user selecting, highlighting, or otherwise identifying a particular point of attention, object, or other item associated with one or more of the devices 105. As such, the current device interaction may provide context for sharpening recognition, interpretation, and understanding of the accompanying utterance, and moreover, the current utterance may provide information to enhance the context provided by the accompanying device interaction.

[028] In one implementation, the voice-click module 108 may determine the various voice-click interactions to be detected based on particular characteristics of the non-voice input devices 105b (e.g., the voice-click interactions may include distinguishable interactions that the non-voice input devices 105b support). For example, a multi-touch display typically includes a

touch-screen display device configured to support various distinguishable gestures for interacting with information displayed therein (e.g., a user may zoom in, zoom out, rotate, or otherwise control graphical information displayed on a multi-touch screen using specific gestures or other interaction techniques). Thus, in one example, the non-voice input devices 105b may include a multi-touch display, in which case the voice-click module 108 may be configured to detect an occurrence of a voice-click when a user engages in one or more of the distinguishable gestures supported by the non-voice multi-touch display 105b.

[029] In one implementation, a user may customize or otherwise modify the voice-click interactions to be detected by the voice-click module 108. In particular, the specific device interactions detected by the voice-click module 108 may be removed or modified, or new device interactions may be added. As such, the voice-click device interactions detected by the voice-click module 108 may include any appropriate interaction or combination of interactions that the non-voice input devices 105b and/or the voice-click module 108 can distinguish.

[030] When the voice-click module 108 detects a user's engaging in a voice-click device interaction, the voice-click module 108 may extract context information associated with the voice-click device interaction for voice-activation. In particular, the voice-click module 108 may identify information relating to a segment, item, point of focus, attention focus, or other data selected by the user, or otherwise identify information relating to a particular device interaction or sequence of device interactions engaged in by the user. The voice-click module 108 thus extracts the identified information relating to the detected voice-click, which may be used as context information to be associated with one or more prior, contemporaneous, or subsequent natural language utterances.

[031] Thus, in response to the voice-click module 108 detecting a voice-click (e.g., selection of an icon, a section of text, specific coordinates on a map display, or other information), the voice-click module 108 may signal the system 100 to use a natural language utterance voice input (which may be received via the voice input device 105a) as further context for determining an action, query, command, task, or other request to perform in order to service the detected voice-click. As such, the various natural language processing components in the system 100 may use the combined context of the voice-click and the accompanying natural language utterance to determine the intent of the voice-click device interaction, and to

appropriately route one or more actions, queries, commands, tasks, or other requests to any of the various devices deployed in the multi-device environment.

[032] For instance, in one implementation, the multi-device environment may include a voice-enabled navigation device. Thus, an exemplary voice-click device interaction may include a user touching a stylus to a specific intersection displayed on a touch-screen display 105b associated with the voice-enabled navigation device, while also providing an utterance such as “What restaurants are around here?” into a microphone 105a. In this example, the voice-click module 108 may extract information relating to the voice-clicked intersection, which may be used as context for processing the accompanying utterance (i.e., the selected intersection may provide context for interpreting “around here,” as opposed to the user’s current location or some other meaning). Moreover, as indicated above, the voice input may be used as additional context for determining task specification. Thus, the utterance may be further processed for recognition and conversational interpretation using the various natural language processing components of the system 100, as described in greater detail below.

[033] In one implementation, an Automatic Speech Recognizer (ASR) 110 may generate one or more preliminary interpretations of the utterance received via the voice input device 105a. For example, the ASR 110 may recognize syllables, words, phrases, or other acoustic characteristics of the utterance using one or more dynamically adaptable recognition grammars. In one implementation, the dynamic recognition grammars may be used to recognize a stream of phonemes using phonetic dictation based on one or more acoustic models (e.g., as described in co-pending U.S. Patent Application 11/197,504, entitled “Systems and Methods for Responding to Natural Language Speech Utterance,” filed August 5, 2005, the disclosure of which is hereby incorporated by reference in its entirety).

[034] In one implementation, the ASR 110 may be configured to perform multi-pass speech recognition, where a primary speech recognition engine may generate a primary transcription of the utterance (e.g., using a large list dictation grammar), and may subsequently request one or more secondary transcriptions from one or more secondary speech recognition engines (e.g., using a virtual dictation grammar having decoy words for out-of-vocabulary words). In one implementation, the primary speech recognition engine may request the secondary transcriptions based on a confidence level for the primary transcription.

[035] The recognition grammars employed in the ASR 110 may include various vocabularies, dictionaries, syllables, words, phrases, or other information for recognizing utterances. In one implementation, information contained in the recognition grammars may be dynamically optimized to improve a likelihood of accurate recognition for a given utterance (e.g., following an incorrect interpretation of a word or phrase, the incorrect interpretation may be removed from the grammar to reduce a likelihood of repeating the incorrect interpretation). Additionally, various forms of knowledge can be used to continually optimize the information contained in the recognition grammars on a dynamic basis. For example, the system 100 may have knowledge that includes environmental knowledge (e.g., peer-to-peer affinities, capabilities of the various devices in the environment, etc.), historical knowledge (e.g., frequent requests, prior context, etc.), or short term shared knowledge relating to a current conversational dialogue or interaction, among other types of knowledge.

[036] In one implementation, the information in the recognition grammars may be further optimized according to context or application-specific domains. In particular, similar utterances may be interpreted differently depending on a context to which the utterance relates, including navigation, music, movies, weather, shopping, news, languages, temporal or geographic proximities, or other contexts or domains. For example, an utterance containing the word “traffic” may be subject to different interpretations depending on whether the context relates to navigation (i.e., road conditions), music (i.e., the 1960’s rock band), or movies (i.e., the film directed by Steven Soderbergh). Accordingly, the ASR 110 may use various techniques to generate a preliminary interpretation of the natural language utterance, such as those described in the above-referenced co-pending U.S. Patent Applications and/or co-pending U.S. Patent Application 11/513,269, entitled “Dynamic Speech Sharpening,” filed August 31, 2006, the disclosure of which is hereby incorporated by reference in its entirety.

[037] As such, the ASR 110 may provide one or more preliminary interpretations of the natural language utterance included in the voice-click to a conversational language processor 120. The conversational language processor 120 may include various natural language processing components collectively configured to model human-to-human conversations or interactions. For example, the conversational language processor 120 may include, among other things, an intent determination engine 130a, a constellation model 130b, one or more domain agents 130c, a context tracking engine 130d, a misrecognition engine 130e, and a voice

search engine 130f. Furthermore, the conversational language processor 120 may be coupled to one or more data repositories 160 and one or more applications 150 associated with various contexts or domains.

[038] Thus, the system 100 may use the various natural language processing components associated with the conversational language processor 120 in order to engage the user in a cooperative conversation and resolve voice-click device interactions based on the user's intent in initiating the voice-click. More particularly, the intent determination engine 130a may establish meaning for a given multi-modal device interaction based on capabilities of the system 100 as well as the capabilities of any other devices in the multi-device environment. For instance, referring to the example above where the user voice-clicked a particular intersection to determine "What restaurants are around here," the conversational language processor 120 may determine a conversational goal of the voice-click (e.g., "What" may indicate that the utterance relates to a query requesting data retrieval). In addition, the conversational language processor 120 may invoke the context tracking engine 130d to determine context for the voice-click. For example, to determine the voice-click context, the context tracking engine 130d may combine the context associated with the identified point of focus (i.e., the selected intersection) with context associated with the utterance (i.e., restaurants).

[039] As a result, the combined context of the voice-click (which includes both the device interaction and the accompanying utterance) may provide sufficient information for routing a specific query. For example, the query may include various parameters or criteria relating to restaurants and the identified intersection. The conversational language processor 120 may then select a particular device, application, or other component to which the query may be routed for processing. For example, in one implementation, the conversational language processor 120 may evaluate the constellation model 130b, which contains a model of capabilities for each device in the multi-device environment. In one implementation, the constellation model 130b may contain, among other things, knowledge of processing and storage resources available to each of the devices in the environment, as well as the nature and scope of domain agents, context, capabilities, content, services, or other information for each of the devices.

[040] As such, using the constellation model 130b and/or other information, the conversational language processor 120 may determine which device or combination of devices

has suitable capabilities that can be invoked to process a given voice-click device interaction. For instance, again referring to the example given above, the conversational language processor 120 may determine that the context of the voice-click relates to an interaction with a navigation device and therefore route the query for processing using a navigation application 150. Results of the query may then be processed (e.g., to weigh the results based on knowledge of the user, such as a preference for vegetarian restaurants) and returned to the user via an output device 180.

[041] According to various aspects of the invention, Figure 2 illustrates a block diagram of an exemplary method for synchronizing multi-modal devices in a natural language voice service environment. As described above, multi-modal device interactions (or “voice-clicks”) may generally occur when a user engages in one or more interactions with one or more multi-modal devices while also providing one or more natural language utterances that relate to the interactions with the multi-modal devices. In one implementation, context information relating to the interactions with the multi-modal devices may be combined with context information relating to the natural language utterances to determine an intent of the voice-click (e.g., to initiate a particular action, query, command, task, or other request).

[042] In one implementation, various natural language processing components may be configured to continually listen or otherwise monitor the multi-modal devices to determine when voice-clicks occur. As such, the method illustrated in Figure 2 may be used to calibrate or otherwise configure the components responsible for continually listening or otherwise monitoring the multi-modal devices. For example, in one implementation, the natural language voice service environment may include a plurality of multi-modal devices that provide different capabilities or services, and the user may engage in one or more voice-clicks to request services relating to any of the various devices or capabilities in any given device interaction.

[043] To enable continuous listening for multi-modal device interactions or voice-clicks, each of the plurality of devices in the environment may be configured to receive information relating to a voice-click. Thus, in one implementation, an operation 210 may include establishing device listeners for each of the plurality of devices in the environment. Additionally, operation 210 may be performed in response to one or more new devices being added to the environment. The device listeners established in operation 210 may include any suitable combination of instructions, firmware, or other routines configured to be executed on one or more processing

devices or other hardware components. For each device in the environment, the associated device listener may communicate with the device to determine capabilities, features, supported domains, or other information associated with the device. In one implementation, the device listeners may be configured to communicate with the device using the Universal Plug and Play protocol designed for ancillary computer devices. It will be apparent, however, that any suitable mechanism for communicating with the multi-modal devices may be used.

[044] When device listeners have been established for each of the devices in the environment (or when device listeners have been established for new devices added to the environment), the various device listeners may be synchronized in an operation 220. In particular, each of the various devices may have different internal clocks or other timing mechanisms, wherein operation 220 may include synchronizing the various device listeners according to the respective internal clocks or timing mechanisms of the devices. In one implementation, synchronizing the device listeners may include each of the respective device listeners publishing information relating to the internal clock or timing of the associated device.

[045] Thus, when one or more multi-modal interactions or voice-clicks subsequently occur for one or more of the devices, the associated device listener may detect information associated with the voice-click in an operation 230. For example, in one implementation, the various device listeners established in operation 210 may be associated with the voice-click module described above and illustrated in Figure 1. Operation 230 may therefore include one or more of the device listeners or the voice-click module detecting an occurrence of a user interacting with one or more of the devices (e.g., selecting data associated with the device, identifying a point of focus or attention focus associated with the device, or otherwise engaging in one or more interactions or sequences of interactions with the device). Furthermore, an operation 240 may then comprise capturing an utterance from the user that relates to the device interaction detected in operation 230.

[046] For example, a user viewing a web page presented on a display device may see a product name on the web page and desire more information about purchasing the product. The user may select text from the web page that contains the product name (e.g., using a mouse or keyboard to highlight the text), and then initiate a voice-click to ask “Is this available on Amazon.com?” In this example, operation 230 may include a device listener associated with the display device detecting the selection of the text associated with the product name, while

operation 240 may include capturing the utterance inquiring about the availability of the product on Amazon.com.

[047] As described above, each device that receives an input from the user may have an internal clock or timing mechanism. In an operation 250, each device may therefore determine when the input was received from a local perspective, and notify the voice-click module that the input was received. In particular, a given voice-click may include at least a natural language utterance received via a voice input device in addition to one or more further interactions with one or more other devices. The utterance may be received prior to, contemporaneously with, or subsequent to the device interactions, whereby operation 250 includes determining the timing of the device interaction for correlation with the associated utterance. In particular, using the device listener signals synchronized as described in reference to operation 220, an operation 260 may include aligning the signals for the device interactions and the utterance. In matching the device interaction and utterance signals, a voice-click input may be generated that includes aligned voice and non-voice components. The voice-click input may then be subject to further natural language processing, as described in greater detail below.

[048] According to various aspects of the invention, Figure 3 illustrates a flow diagram of an exemplary method for processing multi-modal device interactions in a natural language voice service environment. As described above, multi-modal device interactions (or “voice-clicks”) may generally occur when a user interacts with one or more multi-modal devices while also providing one or more natural language utterances related to the device interactions. As such, in one implementation, the method illustrated in Figure 3 may be performed when one or more natural language processing components continually listen or otherwise monitor the one or more multi-modal devices to determine when one or more voice-clicks occur.

[049] In one implementation, one or more device interactions may be defined as initiating a voice-click. For example, any given electronic device may generally support various distinguishable interactions, which may result in a given action, command, query, or other request being executed. Thus, in one implementation, any appropriate combination of device interactions that a given electronic device can uniquely recognize or otherwise use to generate a uniquely recognizable signal may be defined as a voice-click, wherein the voice-click may provide a signal indicating that a natural language utterance is to be processed together with context associated with the relevant device interactions. For example, devices that have a four-

way or a five-way navigation button may support specific distinguishable interactions, wherein pressing the navigation button in a particular way may cause a specific task or other action to be performed, such as controlling a map display or calculating a route. In another example, a BlackBerry® device having a wheel selector may support interactions such as wheeling a cursor over a particular point of focus or attention focus, pressing the wheel to select specific data or a given application, or various other interactions. Various other device interactions may be used to indicate when a natural language utterance is to be processed together with context associated with the device interactions, without limitation, wherein the specific device interactions may vary in any given implementation. For example, the relevant device interactions may further include one or more of gesturing with a pointing or drawing instrument on a touch-sensitive screen (e.g., drawing an ear-shaped squiggle), a unique interaction method such as a long touch or double-tap, and/or if the system is operating in the continuous listening mode described above, a predefined contextual command word may indicate that a current device context is to be processed with a portion of the voice-based input that follows the contextual command word (e.g., the command word be “OK,” “Please,” “Computer,” or another suitable word, wherein a user may select a specific point on a map and say “Please zoom in,” or say “OK read it” when an e-mail is displayed, etc.).

[050] As such, an operation 310 may include processing multi-modal device interactions in a natural language voice service environment to detect an occurrence of one or more device interaction that signal initiation of a voice-click. In particular, the device interactions detected in operation 310 may include any suitable interactions that cause an electronic device to generate a unique, recognizable, or otherwise distinguishable signal that relates to a user selecting data, identifying a point of focus or attention focus, invoking an application or task, or interacting with the device in another way, depending on the specific capabilities of the device.

[051] In addition to the specific signal that the device generates in response to the user interaction, the interaction detected in operation 310 may signal initiation of a voice-click, whereby a prior, contemporaneous, or subsequent natural language voice input will provide further context for interpreting the device interaction detected in operation 310. For example, a natural language processing system may generally be configured to accept a voice input when particular device interactions occur (e.g., pressing a button to turn on a microphone). Thus, in the method illustrated in Figure 3, the device interactions that signal an incoming voice input

may further include any suitable interaction or combination of interactions with an electronic device, including interactions associated with a user selecting data, identifying a point of focus or attention focus, invoking an application or task, or interacting with the device in another way, depending on the specific capabilities of the device.

[052] As such, when a voice-click device interaction has been detected in operation 310, a voice-click signal may be generated in an operation 320 to indicate that a natural language voice input should be associated with the interaction detected in operation 320. Subsequently, an operation 330 may include capturing a user utterance to be associated with the interaction detected in operation 310. In one implementation, the interaction detected in operation 310 may signal that a subsequent voice input will be provided, but it will be apparent that in one or more implementations the utterance captured in operation 330 may be provided prior to or contemporaneously with the interaction detected in operation 310 (e.g., a user may provide an utterance such as “Look up this artist on iTunes®” and subsequently voice-click the artist’s name on a media player, or the user may provide the utterance while voice-clicking the artist’s name, or the user may voice-click the artist’s name and then provide the utterance).

[053] When information relating to the voice-click device interactions and associated natural language utterance have been received, an operation 340 may include extracting and combining context information for the device interactions and the associated utterance. In particular, the context information extracted from the voice-click device interactions may include information relating to a segment, item, point of focus, attention focus, or data selected by the user, or to a particular device interaction or sequence of device interactions engaged in by the user. The extracted context for the device interaction may then be combined with context extracted for the natural language utterance captured in operation 330, wherein the combined context information may be used to determine an intent of the voice-click in an operation 350.

[054] For example, in an exemplary voice-click device interaction, a user may be selectively copying a music collection from a media player to a backup storage device. While browsing the music on the media player, the user may encounter a particular song and voice-click the song while saying “Copy this whole album” (e.g., by pressing a particular button on the media player for an extended period of time while highlighting the song). In this example, operation 310 may include detecting the interaction of the extended button press, which causes the voice-click

signal to be generated in operation 320. The utterance of “Copy this whole album” may then be captured in operation 330, and context information associated with the voice-click device interaction and the utterance may be combined in operation 340. In particular, the context of the device interaction may include information relating to the selected song, among other things (e.g., the context may further include information contained in metadata associated with the song, such as an ID3 tag for music files). Furthermore, the context of the utterance may include information identifying a copy operation and an album that contains the selected song.

[055] As such, context information relating to voice-click interactions with the multi-modal devices may be combined with context information relating to the natural language utterances, whereby an operation 350 may determine the intent of the voice-click interaction. For instance, referring to the example above, the intent determined in operation 350 may include an intent to copy an album that contains the highlighted song from the media player to the backup storage device. Thus, in response to determining the intent of the voice-click in operation 350, one or more requests may be appropriately routed in an operation 360. In the example being discussed herein, operation 360 may include routing one or more requests to the media player to identify all of the data associated with the album that contains the highlighted song, and one or more requests to any appropriate combination of devices that can manage copying of the identified data from the media player to the backup storage device (e.g., a personal computer that interfaces with both the media player and the storage device).

[056] According to various aspects of the invention, Figure 4 illustrates a flow diagram of an exemplary method for processing multi-modal device interactions to generate a transaction lead or “click-through” in a natural language voice service environment. In particular, the method illustrated in Figure 4 may be used to generate a transaction lead or click-through in combination with one or more actions performed in response to one or more voice-click device interactions being detected.

[057] For example, an operation 410 may include detecting one or more voice-click device interactions received from a user, wherein the voice-click device interactions may comprise any suitable combination of one or more device interactions coupled with one or more related natural language utterances. The user’s intent in engaging in the voice-click device interactions may then be determined in an operation 420, and a subsequent operation 430 may include routing one or more requests to one or more processing devices based on the determined

intent in order to resolve the voice-click interaction. In one implementation, operations 410, 420, and 430 may be performed in a manner similar to that described above in reference to Figure 2 and Figure 3, whereby signals for the device interactions may be aligned with signals for one or more natural language utterances and context information may be extracted from the signals to determine the intent of the voice-click device interaction.

[058] In addition to routing one or more requests based on the user's intent, the method illustrated in Figure 4 may further comprise generating one or more transaction leads that may result in one or more click-throughs. For example, a click-through may generally refer to an instance of a user clicking or otherwise selecting an electronic advertisement to access one or more services associated with the advertiser. In many electronic systems, click-throughs or click-through rates can provide mechanisms for measuring a user's interactions with an electronic advertisement, which may provide various measurements that advertisers can use to determine an amount to pay an entity that delivers the advertisement to users.

[059] As such, the method illustrated in Figure 4 may generate transaction leads, including advertisements or recommendations, whereby a user's voice-based input combined with a specific device interaction may provide further focus for generating transaction leads. In this manner, the advertisements or recommendations provided to the user may be more relevant to specific information with which the user may be interacting. Furthermore, using natural language cognitive models and shared knowledge relating to a user's preferences may provide further context for targeted transaction leads tailored to the specific user, which may thus be more likely to result in a click-through that can generate payment for a voice services provider.

[060] Thus, in addition to routing one or more requests based on the user's intent in engaging in the voice-click device interaction, an operation 440 may include generating one or more transaction leads based on the determined intent. In particular, based on the combined context of the device interaction and the associated natural language utterance, the transaction leads may be processed in a manner "closer" to the user, in that local voice and non-voice context can be used as state data in any appropriate system that performs targeted advertising. For instance, referring to the example given above where a user selects an intersection displayed on a navigation device while saying "Find restaurants are around here," the transaction leads generated in operation 440 may include one or more advertisements or recommendations for restaurants near the intersection, which may be targeted to the user

based on knowledge of the user's short-term and long-term preferences (e.g., a preferred type of restaurant, a preferred price range, etc.).

[061] The transaction leads may then be presented to the user (e.g., as selectable points on a map display). The user's subsequent multi-modal device interactions may then be monitored in an operation 450 to determine whether or when one or more further multi-modal device interactions occur. If no further interactions occur, a determination may be made that the user did not act on the transaction lead and the process may then end. If an additional multi-modal interaction does occur, however, the multi-modal input may be processed in an operation 480 to determine an intent of the input and route one or more requests accordingly. In addition, an operation 460 may include determining whether the multi-modal input relates to the transaction lead generated in operation 440. For example, the user may select one of the advertised or recommended restaurants by providing an utterance, a non-voice device interaction, or a voice-click device interaction requesting further action or information relating to the transaction lead. In such a case, an operation 470 may further include processing a transaction click-through in relation to the transaction lead generated in operation 440, wherein the transaction click-through may be used to determine payment or otherwise generate revenue for a particular entity (e.g., a provider of the voice services or another entity associated with the transaction lead or transaction click-through).

[062] Implementations of the invention may be made in hardware, firmware, software, or various combinations thereof. The invention may also be implemented as instructions stored on a machine-readable medium, which may be read and executed by one or more processors. A machine-readable medium may include various mechanisms for storing or transmitting information in a form readable by a machine (e.g., a computing device). For example, a machine-readable storage medium may include read only memory, random access memory, magnetic disk storage media, optical storage media, flash memory devices, or other storage media, and a machine-readable transmission media may include forms of propagated signals, such as carrier waves, infrared signals, digital signals, or other transmission media. Further, firmware, software, routines, or instructions may be described in the above disclosure in terms of specific exemplary aspects and implementations of the invention, and performing certain actions. However, it will be apparent that such descriptions are merely for convenience, and

that such actions in fact result from computing devices, processors, controllers, or other devices executing the firmware, software, routines, or instructions.

[063] Although the descriptions provided herein have generally focused on techniques for processing multi-modal device interactions in a natural language voice services environment, it will be apparent that various further natural language processing capabilities may be used in connection with, in addition to, or in lieu of those described in connection with the specific aspects and implementations discussed herein. For example, in addition to the techniques described in the co-pending U.S. Patent Applications referenced above, the systems and methods described herein may further utilize natural language processing capabilities as described in co-pending U.S. Patent Application 11/197,504, entitled “Systems and Methods for Responding to Natural Language Speech Utterance,” filed August 5, 2005, U.S. Patent Application 11/200,164, entitled “System and Method of Supporting Adaptive Misrecognition in Conversational Speech,” filed August 10, 2005, U.S. Patent Application 11/212,693, entitled “Mobile Systems and Methods of Supporting Natural Language Human-Machine Interactions,” filed August 29, 2005, U.S. Patent Application 11/580,926, entitled “System and Method for a Cooperative Conversational Voice User Interface,” filed October 16, 2006, U.S. Patent Application 11/671,526, entitled “System and Method for Selecting and Presenting Advertisements Based on Natural Language Processing of Voice-Based Input,” filed February 6, 2007, and U.S. Patent Application 11/954,064, entitled “System and Method for Providing a Natural Language Voice User Interface in an Integrated Voice Navigation Services Environment,” filed December 11, 2007, the disclosures of which are hereby incorporated by reference in their entirety.

[064] Accordingly, aspects and implementations of the invention may be described herein as including a particular feature, structure, or characteristic, but it will be apparent that every aspect or implementation may or may not necessarily include the particular feature, structure, or characteristic. In addition, when a particular feature, structure, or characteristic has been described in connection with a given aspect or implementation, it will be understood that such feature, structure, or characteristic may be included in connection with other aspects or implementations, whether or not explicitly described. Thus, various changes and modifications may be made to the preceding description without departing from the scope or spirit of the

invention, and the specification and drawings should therefore be regarded as exemplary only, with the scope of the invention determined solely by the appended claims.

CLAIMS

What is claimed is:

1. A method for processing one or more multi-modal device interactions in a natural language voice services environment that includes one or more electronic devices, comprising:
 - detecting at least one multi-modal device interaction, wherein the multi-modal device interaction includes a non-voice interaction with at least one of the electronic devices or an application associated with at least one of the electronic devices, and wherein the multi-modal device interaction further includes at least one natural language utterance relating to the non-voice interaction;
 - extracting context information relating to the multi-modal device interaction, wherein the extracted context information includes context relating to the non-voice interaction, and wherein the extracted context information further include context relating to the natural language utterance;
 - combining the context relating to the non-voice interaction and the context relating to the natural language utterance;
 - determining an intent of the multi-modal device interaction based on the combined context relating to the non-voice interaction and the natural language utterance; and
 - routing at least one request to one or more of the electronic devices based on the determined intent of the multi-modal device interaction.
2. The method of claim 1, wherein at least one of the electronic devices includes an input device configured to receive the natural language utterance.
3. The method of claim 2, further comprising signaling the input device to capture the natural language utterance in response to the non-voice interaction being detected.
4. The method of claim 3, further comprising:
 - establishing one or more device listeners in the natural language voice services environment, the device listeners configured to detect the non-voice interaction; and
 - aligning information relating to the non-voice interaction detected by the device listeners and the natural language utterance captured by the input device.

5. The method of claim 1, further comprising:
 - generating at least one transaction lead based on the determined intent of the multi-modal device interaction;
 - receiving at least one further multi-modal device interaction relating to the generated transaction lead; and
 - processing a transaction click-through in response to receiving the multi-modal device interaction relating to the generated transaction lead.
6. The method of claim 5, wherein the generated transaction lead includes at least one of an advertisement or a recommendation relating to the determined intent of the multi-modal device interaction.
7. The method of claim 1, wherein the non-voice interaction comprises a selection of a segment, item, data, or application associated with one or more of the electronic devices.
8. The method of claim 1, wherein the non-voice interaction comprises an identification of a point of focus or an attention focus associated with one or more of the electronic devices.
9. The method of claim 1, wherein the non-voice interaction comprises one or more unique and distinguishable interactions associated with one or more of the electronic devices.
10. A system for processing one or more multi-modal device interactions in a natural language voice services environment that includes one or more electronic devices, wherein the system comprises one or more processing devices configured to:
 - detect at least one multi-modal device interaction, wherein the multi-modal device interaction includes a non-voice interaction with at least one of the electronic devices or an application associated with at least one of the electronic devices, and wherein the multi-modal device interaction further includes at least one natural language utterance relating to the non-voice interaction;

extract context information relating to the multi-modal device interaction, wherein the extracted context information includes context relating to the non-voice interaction, and wherein the extracted context information further include context relating to the natural language utterance;

combine the context relating to the non-voice interaction and the context relating to the natural language utterance;

determine an intent of the multi-modal device interaction based on the combined context relating to the non-voice interaction and the natural language utterance; and

route at least one request to one or more of the electronic devices based on the determined intent of the multi-modal device interaction.

11. The system of claim 10, wherein at least one of the electronic devices includes an input device configured to receive the natural language utterance.

12. The system of claim 11, the processing devices further configured to signal the input device to capture the natural language utterance in response to the non-voice interaction being detected.

13. The system of claim 12, the processing devices further configured to:
establish one or more device listeners in the natural language voice services environment, the device listeners configured to detect the non-voice interaction; and
align information relating to the non-voice interaction detected by the device listeners and the natural language utterance captured by the input device.

14. The system of claim 10, the processing devices further configured to:
at least one transaction lead based on the determined intent of the multi-modal device interaction;
receive at least one further multi-modal device interaction relating to the generated transaction lead; and
process a transaction click-through in response to receiving the multi-modal device interaction relating to the generated transaction lead.

15. The system of claim 14, wherein the generated transaction lead includes at least one of an advertisement or a recommendation relating to the determined intent of the multi-modal device interaction.
16. The system of claim 10, wherein the non-voice interaction comprises a selection of a segment, item, data, or application associated with one or more of the electronic devices.
17. The system of claim 10, wherein the non-voice interaction comprises an identification of a point of focus or an attention focus associated with one or more of the electronic devices
18. The system of claim 10, wherein the non-voice interaction comprises one or more unique and distinguishable interactions associated with one or more of the electronic devices.

1 / 4

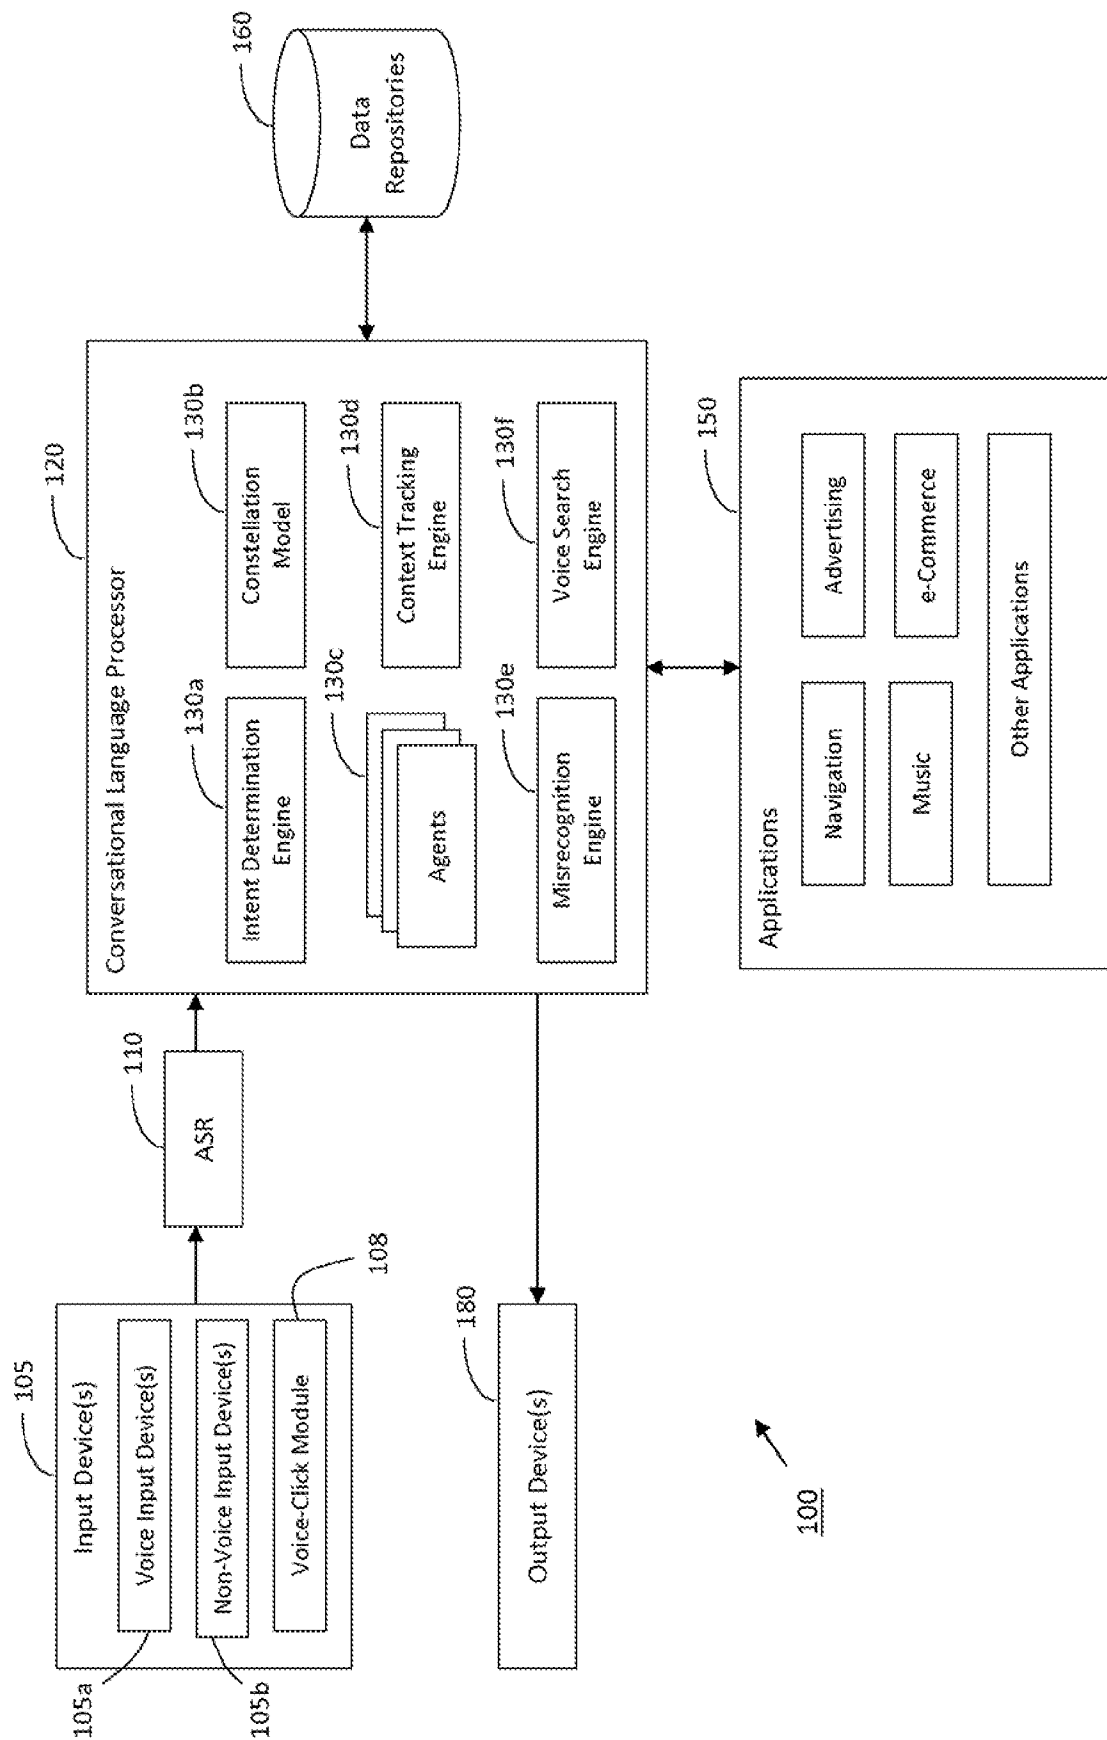


Figure 1

2 / 4

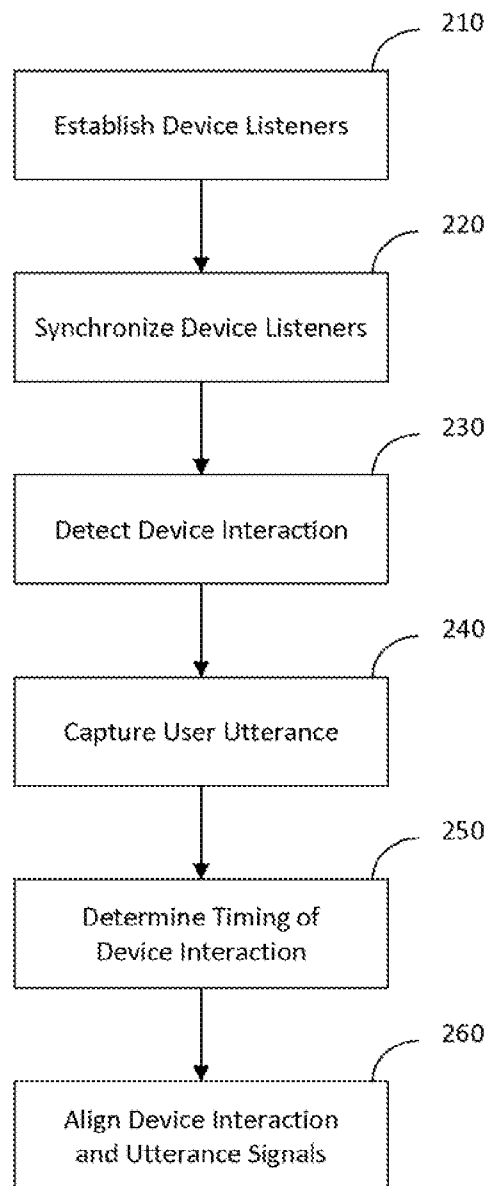


Figure 2

3 / 4

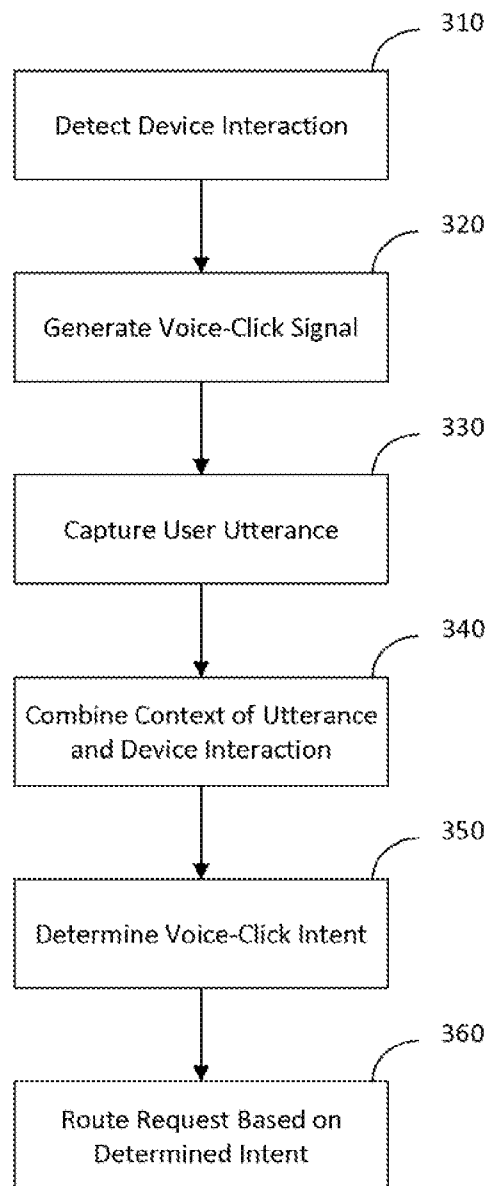


Figure 3

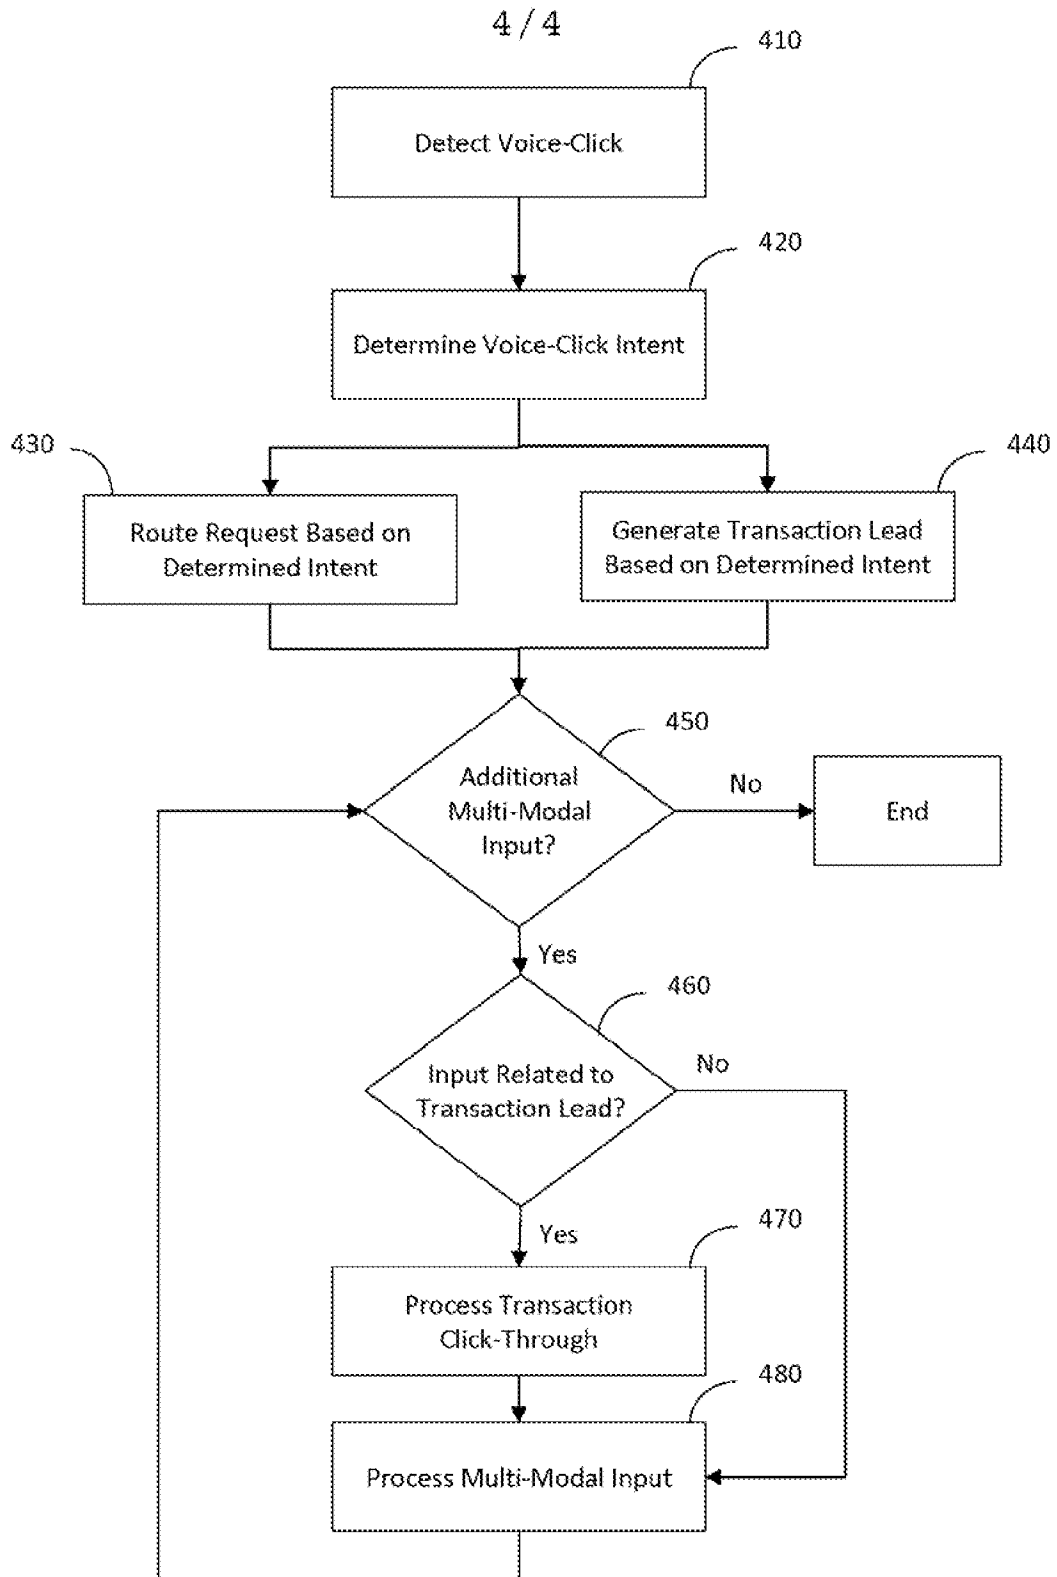


Figure 4

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 10/24895

A. CLASSIFICATION OF SUBJECT MATTER

IPC(8) - G10L 15/00 (2010.01)

USPC - 704/257

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

USPC: 704/257

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

USPC: 704/270.1; 704/200; 704/9 (search term limited)

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Electronic Databases Searched: pubWEST(PGPB,USPT,USOC,EPAB,JPAB); googleScholar

Search Terms Used: natural language, multimodal, speech/voice identification/recognition, voice/keyboard/mouse click, context/intent, marketing/advertising

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X --- Y	US 2007/0050191 A1 (WEIDER et al.) 01 March 2007 (01.03.2007) Entire document, especially: para [0001], [0016], [0018], [0020], [0031], [0043], [0116], [0133]; Fig. 7, 8	1-2, 7-11 and 16-18 ----- 3-6 and 12-15
Y	US 2007/0043574 A1 (COFFMAN et al.) 22 February 2007 (22.02.2007) Entire document, especially: Abstract, para [0154], [0227]	3, 4, 12 and 13
Y	US 2008/0189110 A1 (FREEMAN et al.) 07 August 2008 (07.08.2008) Entire document, especially: para [0006], [0013], [0024], [0028] and Fig. 3	5, 6, 14 and 15
A	US 2008/0235027 A1 (CROSS) 25 September 2008 (25.09.2008)	1-18
A	US 2008/0091406 A1 (BALDWIN et al.) 17 April 2008 (14.04.2008)	1-18

☐ Further documents are listed in the continuation of Box C.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

28 March 2010 (28.03.2010)

Date of mailing of the international search report

13 APR 2010

Name and mailing address of the ISA/US

Mail Stop PCT, Attn: ISA/US, Commissioner for Patents
P.O. Box 1450, Alexandria, Virginia 22313-1450

Facsimile No. 571-273-3201

Authorized officer:

Lee W. Young

PCT Helpdesk: 571-272-4300

PCT OSP: 571-272-7774