



- (51) **International Patent Classification:**
G10L 15/08 (2006.01) *G10L 15/04* (2013.01)
G10L 25/87 (2013.01)
- (21) **International Application Number:**
PCT/US2016/033736
- (22) **International Filing Date:**
23 May 2016 (23.05.2016)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
14/721,252 26 May 2015 (26.05.2015) US
- (71) **Applicant:** NUANCE COMMUNICATIONS, INC. [US/US]; One Wayside Road, Burlington, MA 01803 (US).
- (72) **Inventor:** FANTY, Mark; 108 Union Street, Norfolk, MA 02056 (US).
- (74) **Agent:** WEHNER, Daniel, T.; Wolf, Greenfield & Sacks, P.C., 600 Atlantic Avenue, Boston, MA 02210-2206 (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, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) **Title:** METHODS AND APPARATUS FOR REDUCING LATENCY IN SPEECH RECOGNITION APPLICATIONS




FIG. 2

(57) **Abstract:** The method comprises receive first audio comprising speech from a user of a computing device, detecting an end of speech in the first audio, generating an ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech, determining whether a valid action can be performed by a speech-enabled application installed on the computing device using the ASR result, and processing second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the ASR result.



SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, **Published:**
GW, KM, ML, MR, NE, SN, TD, TG).

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

METHODS AND APPARATUS FOR REDUCING LATENCY IN SPEECH RECOGNITION APPLICATIONS

BACKGROUND

[0001] Some electronic devices, such as smartphones and tablet computers, include or are associated with speech recognition capabilities that enable users to access functionality of the device via speech input. Speech input is processed by an automatic speech recognition (ASR) system, which converts the input audio to recognized text. Electronic devices may also include or be associated with a natural language understanding (NLU) engine that interprets user input and takes an action based upon determined semantic content of the user's input (e.g., by facilitating actions with one or more applications accessible via the electronic device). Virtual agents or virtual assistants are one such class of applications that benefit from NLU processing to assist users in performing functions such as searching for content on a network (e.g., the Internet) and interfacing with other applications. Users can interact with a virtual agent by typing, touch, speech, or some other interface. To determine a meaning of a user input, the NLU engine interprets the user input, and a virtual agent may attempt to infer an action the user wants to perform based on the NLU result.

SUMMARY

[0002] Some embodiments are directed to a computing device including a speech-enabled application installed thereon. The computing device comprises an input interface configured to receive first audio comprising speech from a user of the computing device, an automatic speech recognition (ASR) engine configured to detect based, at least in part, on a threshold time for endpointing, an end of speech in the first audio, and generate a first ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech. The computing device further comprises at least one processor programmed to determine whether a valid action can be performed by the speech-enabled application using the first ASR result, and instruct the ASR engine to process second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the first ASR result.

[0003] Other embodiments are directed to a method. The method comprises receiving, by an input interface of a computing device, first audio comprising speech from a user of the computing device, detecting, by an automatic speech recognition (ASR) engine of the computing device, an end of speech in the first audio, generating, by the ASR engine, an ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech, determining whether a valid action can be performed by a speech-enabled application installed on the computing device using the ASR result, and instructing the ASR engine to process second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the ASR result.

[0004] Other embodiments are directed to a computer-readable storage medium encoded with a plurality of instructions that, when executed by a computing device, performs a method. The method comprises receiving first audio comprising speech from a user of the computing device, detecting an end of speech in the first audio, generating an ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech, determining whether a valid action can be performed by a speech-enabled application installed on the computing device using the ASR result, and processing second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the ASR result.

[0005] It should be appreciated that all combinations of the foregoing concepts and additional concepts discussed in greater detail below (provided that such concepts are not mutually inconsistent) are contemplated as being part of the inventive subject matter disclosed herein.

BRIEF DESCRIPTION OF DRAWINGS

[0006] The accompanying drawings are not intended to be drawn to scale. In the drawings, each identical or nearly identical component that is illustrated in various figures is represented by a like numeral. For purposes of clarity, not every component may be labeled in every drawing. In the drawings:

[0007] FIG. 1 is a schematic illustration of a client-server system for use with speech-enabled applications in accordance with some embodiments;

[0008] FIG. 2 illustrates a process for determining whether a valid action for a speech-enabled application may be performed based on an ASR result in accordance with some embodiments;

[0009] FIG. 3 illustrates a process for reducing premature endpointing in accordance with some embodiments;

[0010] FIG. 4 illustrates a process for dynamically setting a timeout value for an ASR process in accordance with some embodiments;

[0011] FIG. 5 illustrates a process for storing NLU results on local storage of a client device in accordance with some embodiments;

[0012] FIG. 6 illustrates a process for using NLU results stored on local storage of a client device in accordance with some embodiments; and

[0013] FIG. 7 illustrates a process for displaying dynamically-generated hints on a user interface of a client device in accordance with some embodiments.

DETAILED DESCRIPTION

[0014] Users of electronic devices that include or are associated with automatic speech recognition (ASR) and/or natural language understanding (NLU) processing often report that the latency of receiving results after providing speech input is a significant factor influencing a positive user experience. To reduce latency, most or all ASR and/or NLU processing can be performed locally on the device. However, this approach may be undesirable for some devices with limited memory and/or processing resources. As discussed in further detail below, distributed systems where at least a portion of the ASR and/or NLU processing is provided by one or more servers connected to the device (e.g., via one or more networks) in a client/server architecture are frequently used to reduce the resource burden on client devices. Sending and receiving information over network(s) increases the latency associated with ASR and/or NLU processing, and some embodiments discussed further below are directed to techniques for reducing latencies associated with distributed ASR and/or NLU processing to improve the user experience with applications on the client device that use ASR and/or NLU processing.

[0015] Systems that process audio including speech input for ASR typically determine when the user has finished speaking based on an analysis of the speech input.

The process for determining when a user has finished speaking is often referred to as “endpoint detection.” Endpoint detection may be accomplished by determining when the user’s speech has ended for an amount of time that exceeds a threshold or timeout value (e.g., three seconds). The audio including speech is processed by an ASR engine to generate a textual ASR result. The ASR result may then be processed by an NLU engine, as discussed above, to infer an intended action that the user would like to perform.

[0016] Premature endpoint detection (i.e., determining that the user was done speaking before the user has finished the input the user desired to provide) typically results in the processing by an NLU engine of an ASR result with insufficient information for the NLU engine to properly infer an intended action for the user’s desired input. The inventors have recognized and appreciated that premature endpoint detection often arises in cases where a user pauses in the middle of an utterance while the user thinks about what to say next. For example, a user asking a speech-enabled application for directions may say, “I would like directions to,” followed by a pause while the user thinks about the destination location. An ASR result corresponding to the utterance “I would like directions to” would typically be interpreted by an NLU engine associated with the speech-enabled application as either being an error and/or as being incomplete. To enable the NLU engine to determine an NLU result that could be used by a speech-enabled application to perform a valid action, the user may be required to start over or provide additional information in a new utterance, either of which is time consuming and detracts from the user experience.

[0017] A technique for reducing the frequency of premature endpoint detections is to increase the timeout value used for endpoint detection. However, doing so causes a higher latency for all utterances, resulting in an unfavorable user experience. Some embodiments described herein are directed to techniques for selectively and dynamically controlling the endpoint detection process during processing of speech input in an effort to reduce the number of ASR results with insufficient information sent to and processed by the NLU engine.

[0018] Other embodiments are directed to reducing latency associated with remote NLU processing by storing at least some NLU results (e.g., NLU results for

recent, frequent, and/or any other desired type of utterances) in a local storage device of a client device. By locally storing or “caching” at least some NLU results on a client device, the cached NLU results may be obtained with reduced latency compared to server-based NLU processing, and may be obtained even when a network connection to the server is not available.

[0019] Other embodiments are directed to training users on options for interacting with a speech-enabled application (e.g., a voice-enabled virtual assistant application) installed on a client device by displaying dynamically-generated hints on a user interface of the client device. Because the dynamically-generated hints are generated based, at least in part, what the user has already said, the user learns how to interact with the speech-enabled application in future utterances and the hint also provides the user with information regarding how to complete an input in progress.

[0020] The techniques described herein may be implemented in any application or system that uses NLU-based processing. In some embodiments, described below, an NLU system using the techniques described herein may be used to facilitate interactions between a user and a virtual assistant (e.g., implemented as an application executing on an electronic device such as a smartphone). However, this is but one illustrative use for the techniques described herein, as they may be used with any NLU system in any environment.

[0021] FIG. 1 shows an illustrative computing environment 100 that may be used in accordance with some embodiments of the invention. Computing environment 100 includes electronic device 110. In some embodiments, electronic device 110 may be a client device in a client-server architecture, as discussed in more detail below. Electronic device 110 includes input interface 112 configured to receive user input. The input interface may take any form as the aspects of the invention are not limited in this respect. In some embodiments, input interface 112 may include multiple input interfaces each configured to receive one or more types of user input. For example, input interface 112 may include a keyboard (e.g., a QWERTY keyboard), a keypad, a touch-sensitive screen, a mouse, or any other suitable user input device. As another example, input interface may include a microphone that, when activated, receives speech input, and the system may perform automatic speech recognition (ASR) either locally on the electronic device,

remotely (e.g., on a server), or distributed between both. The received speech input may be stored in a datastore (e.g., local storage 140) associated with electronic device 110 to facilitate the ASR processing.

[0022] Electronic device 110 also includes output interface 114 configured to output information from electronic device 110. The output interface may take any form, as aspects of the invention are not limited in this respect. In some embodiments, output interface 114 may include multiple output interfaces each configured to provide one or more types of output. For example, output interface 114 may include one or more displays, one or more speakers, or any other suitable output device. Applications executing on electronic device 110 may be programmed to display a user interface to facilitate the performance of one or more actions associated with the application. In one example described herein, the application displays a user interface that provides dynamically-generated hints to a user to help the user complete an input in progress, as described in more detail below.

[0023] Electronic device 110 also includes one or more processors 116 programmed to execute a plurality of instructions to perform one or more functions on electronic device 110. Exemplary functions include, but are not limited to, facilitating the storage of user input, launching and executing one or more applications on electronic device 110, and providing output information via output interface 114. Exemplary functions also include performing speech recognition (e.g., using ASR engine 130) and performing natural language understanding (e.g., using NLU system 132), as discussed in more detail below.

[0024] Electronic device 110 also includes network interface 122 configured to enable electronic device 110 to communicate with one or more computers via network 120. Some embodiments may be implemented using a client/server system where at least a portion of an ASR and/or an NLU process is performed remotely from electronic device 110. In such embodiments, network interface 122 may be configured to provide information to one or more server devices 150 to perform ASR, an NLU process, both ASR and an NLU process, or some other suitable function. Server 150 may be associated with one or more non-transitory datastores (e.g., remote storage 160) that facilitate processing by the server.

[0025] Network 120 may be implemented in any suitable way using any suitable communication channel(s) enabling communication between the electronic device and the one or more computers. For example, network 120 may include, but is not limited to, a local area network, a wide area network, an Intranet, the Internet, wired and/or wireless networks, or any suitable combination of local and wide area networks. Additionally, network interface 122 may be configured to support any of the one or more types of networks that enable communication with the one or more computers.

[0026] In some embodiments, electronic device 110 is configured to process speech received via input interface 112, and to produce at least one speech recognition result using ASR engine 130. ASR engine 130 is configured to process audio including speech using automatic speech recognition to determine a textual representation corresponding to at least a portion of the speech. ASR engine 130 may implement any type of automatic speech recognition to process speech, as the techniques described herein are not limited to the particular automatic speech recognition process(es) used. As one non-limiting example, ASR engine 130 may employ one or more acoustic models and/or language models to map speech data to a textual representation. These models may be speaker independent or one or both of the models may be associated with a particular speaker or class of speakers. Additionally, the language model(s) may include domain-independent models used by ASR engine 130 in determining a recognition result and/or models that are tailored to a specific domain. The language model(s) may optionally be used in connection with a natural language understanding (NLU) system (e.g., NLU system 132), as discussed in more detail below. ASR engine 130 may output any suitable number of recognition results, as aspects of the invention are not limited in this respect. In some embodiments, ASR engine 130 may be configured to output N-best results determined based on an analysis of the input speech using acoustic and/or language models, as described above.

[0027] Electronic device 110 also includes NLU system 132 configured to process a textual representation to gain some semantic understanding of the input, and output one or more NLU hypotheses based, at least in part, on the textual representation. In some embodiments, the textual representation processed by NLU system 132 may comprise one or more ASR results (e.g., the N-best results) output from an ASR engine (e.g., ASR engine 130), and the NLU system may be configured to generate one or more

NLU hypotheses for each of the ASR results. It should be appreciated that in addition to an ASR result, NLU system 132 may also process other suitable textual representations. For example, a textual representation entered via a keyboard, a touch screen, or received using some other input interface may additionally be processed by an NLU system in accordance with the techniques described herein. Additionally, text-based results returned from a search engine or provided to electronic device 110 in some other way may also be processed by an NLU system in accordance with one or more of the techniques described herein. The NLU system and the form of its outputs may take any of numerous forms, as the techniques described herein are not limited to use with NLU systems that operate in any particular manner.

[0028] The electronic device 110 shown in FIG. 1 includes both ASR and NLU processes being performed locally on the electronic device 110. In some embodiments, one or both of these processes may be performed in whole or in part by one or more computers (e.g., server 150) remotely located from electronic device 110. For example, in some embodiments that include an ASR component, speech recognition may be performed locally using an embedded ASR engine associated with electronic device 110, a remote ASR in network communication with electronic device 110 via one or more networks, or speech recognition may be performed using a distributed ASR system including both embedded and remote components. In some embodiments, NLU system 132 may be located remotely from electronic device 110 and may be implemented using one or more of the same or different remote computers configured to provide some or all of the remote ASR processing. Additionally, it should be appreciated that computing resources used in accordance with any one or more of ASR engine 130 and NLU system 132 may also be located remotely from electronic device 110 to facilitate the ASR and/or NLU processes described herein, as aspects of the invention related to these processes are not limited in any way based on the particular implementation or arrangement of these components within a computing environment 100.

[0029] FIG. 2 illustrates a process for assessing whether input speech includes sufficient information to allow a speech-enabled application to perform a valid action, in accordance with some embodiments. In act 210, audio including speech is received. The audio may be received in any suitable way. For example, an electronic device may include a speech input interface configured to receive audio including speech, as

discussed above. The process then proceeds to act 212, where at least a portion of the received audio is processed by an ASR engine to generate an ASR result. As discussed above, any suitable ASR process may be used to recognize at least a portion of the received audio, as aspects of the invention are not limited in this respect. For example, in some embodiments, the ASR engine processes the received audio to detect an end of speech in the audio, and the portion of the audio prior to the detected end of speech is processed to generate the ASR result.

[0030] The process then proceeds to act 214, where it is determined whether the speech in the received audio includes sufficient information to allow a speech-enabled application to perform a valid action. Determining whether the speech includes sufficient information to allow a speech-enabled application to perform a valid action may be determined in any suitable way. In some embodiments, the ASR result output from the ASR engine may be processed by an NLU system to generate an NLU result, and the determination of whether the speech includes sufficient information to allow a speech-enabled application to perform a valid action may be based, at least in part, on the NLU result. As discussed above, some NLU systems may process input text and return an error and/or an indication that the input text is insufficient to enable an application to perform a valid action. The error and/or indication that the text input to the NLU includes insufficient information to allow a speech-enabled application to perform a valid action may be used, at least in part, to determine that the endpointing by the ASR engine was performed prematurely (i.e., the user had not finished speaking a desired input prior to the endpointing process being completed).

[0031] In some embodiments, determining whether the received speech includes sufficient information to allow a speech-enabled application to perform a valid action may be determined based, at least in part, on the content of the ASR result output from the ASR engine. For example, as discussed in more detail below, some embodiments compare the ASR result to one or more prefixes stored in local storage on the electronic device, and the determination of whether the received speech includes sufficient information to allow a speech-enabled application to perform a valid action is made, based, at least in part, on whether the ASR result matches one or more of the locally-stored prefixes.

[0032] If it is determined in act 214 that the speech in the received audio includes sufficient information to allow a speech-enabled application to perform a valid action, the process ends and the utterance is processed as it otherwise would be in the absence of the techniques described herein. Conversely, if it is determined that the speech in the received audio does not include sufficient information to allow a speech-enabled application to perform a valid action, the process proceeds to act 216, where the ASR engine is instructed to process additional audio. The additional audio processed by the ASR engine may be used to supplement the received audio with insufficient information to allow a speech-enabled application to perform a valid action, as discussed in more detail below.

[0033] As discussed above, a factor that may contribute to increased processing latencies during use of a speech-driven application that incorporates NLU techniques is premature endpoint detection. Endpoint detection is typically accomplished by assuming that a user has finished speaking by detecting the end of speech (e.g., by detecting silence) in the received audio, and determining that a particular threshold amount of time (e.g., three seconds) has passed since the end of speech was detected. The timeout value used for determining when to endpoint may be a fixed value set by an application programmer, or may be variable based on the use of one or more speech-enabled applications by user(s) of the electronic device. As discussed above, selecting a long timeout value results in processing latency delays for all utterances, and selecting a short timeout value may result in premature endpoint detection for some utterances (e.g., utterances in which the speaker pauses while thinking about what to say next).

[0034] Some embodiments are directed to techniques for recovering from and/or preventing premature endpoint detection by processing audio after the timeout value used for the endpoint detection process has expired. FIG. 3 illustrates a process for recovering from premature endpoint detection in accordance with some embodiments. In act 310, the end of speech in received audio is determined. The end of speech may be determined in any suitable way, as aspects of the invention are not limited in this respect. For example, the end of speech may be determined based, at least in part, on a detected energy level in the audio signal, or the end of speech may be determined by analyzing one or more other characteristics of the audio signal.

[0035] As discussed above, some conventional endpoint detection techniques determine that a user has stopped speaking after a threshold amount of time has passed following detection of the end of speech. A determination of whether the threshold amount of time has passed may be made in any suitable way. For example, at a time corresponding to when the end of speech is first detected, a timer may be started, and if a particular amount of time elapses after the timer is started it may be determined that the threshold amount of time has passed, and that the speaker had finished speaking.

[0036] In response to detecting the end of speech in the first audio, the process proceeds to act 312, where ASR is performed on the first audio (e.g., the audio received before the detected endpoint) to generate a first ASR result for the first audio. Although performing ASR on the first audio is shown as being performed only after an end of speech is detected in the first audio, in some embodiments, an ASR process may be initiated any time after at least a portion of the first audio is received, as aspects of the invention are not limited in this respect. For example, ASR may be performed on one or more time segments of the first audio prior to detecting the end of speech in the first audio to determine whether the first audio includes a locally-stored prefix, as discussed in more detail below.

[0037] After an ASR result for the first audio has been generated in act 312, the process proceeds to act 314, where NLU is performed on the ASR result to generate a first NLU result. NLU may be performed on all or a portion of the ASR result generated for the first audio and the NLU process may be initiated at any time after at least a portion of the first audio has been recognized using an ASR process. Optionally, the first ASR result may be stored for possible combination with a second ASR result generated based, at least in part, on second received audio, as described in more detail below.

[0038] After the first NLU result is generated in act 314, the NLU result may be analyzed in act 316 to determine whether the first NLU result includes sufficient information to allow a speech-enabled application to perform a valid action or whether additional audio input is necessary to allow a speech-enabled application to perform a valid action. Whether the first NLU result includes sufficient information to allow a speech-enabled application to perform a valid action may be determined in any suitable way. For example, if the speech-enabled application is a navigation application, the

NLU result may be considered to include sufficient information when the result includes both an action to be performed (e.g., “provide directions to”) and information used in performing the action (e.g., a destination).

[0039] If it is determined that the first NLU result includes sufficient information to allow a speech-enabled application to perform a valid action, the NLU result is provided to the speech-enabled application, and the process ends. Otherwise, the process proceeds to act 320, where second audio, including audio recorded after the end of speech is detected in the first audio, is processed. In some embodiments, analyzing the second audio comprises processing at least some of the second audio as part of a single utterance including at least some of the first audio.

[0040] Although shown in FIG. 3 as being initiated upon determining that an NLU result includes insufficient information to allow a speech-enabled application to perform a valid action, in some embodiments, processing the second audio may be initiated prior to determining that the NLU result for the first audio includes such insufficient information. For example, processing of at least some of the second audio may be started immediately following (or shortly thereafter) processing of the first audio so that the time between processing the first audio and processing the second audio is short. In some embodiments, processing of the second audio may additionally or alternatively be initiated by any suitable trigger including, but not limited to, detection of evidence that the user has resumed speaking. In some embodiments, a combination of events may trigger the processing of the second audio. Regardless of the event(s) that triggers the processing of the second audio, information in the second audio may supplement the information in the first audio to reduce premature endpointing for speech-enabled applications, as discussed in more detail below.

[0041] The second audio may be of any suitable duration. In some embodiments, the second audio may comprise audio for a fixed amount of time (e.g., three seconds), whereas in other embodiments, the second audio may comprise audio for a variable amount of time (e.g., based on detecting an end of speech in the second audio). The process then proceeds to act 322, where it is determined whether the second audio includes speech. The determination of whether the second audio includes speech may be

made in any suitable way, for example, using well-known techniques for detecting speech in an audio recording.

[0042] If it is determined in act 322 that the second audio does not include speech, the process proceeds to act 324, where the second audio is discarded and the process ends. If it is determined that the second audio includes speech, the process proceeds to act 326, where ASR is performed on at least a portion of the second audio to generate a second ASR result. In some embodiments, the second ASR result is generated based, at least in part, on an analysis of at least a portion of the first audio and at least a portion of the second audio. In other embodiments, the second ASR result is generated based only on an analysis of at least a portion of the second audio. ASR may be performed on at least a portion of the second audio at any suitable time, and embodiments are not limited in this respect. For example, ASR may be performed any time following detection of speech in the second audio so that the ASR processing may begin before the entire second audio is received.

[0043] The process then proceeds to act 328, where NLU is performed based, at least in part, on the first ASR result and/or the second ASR result. In some embodiments, NLU may be performed based only on the second ASR result to generate a second NLU result, and the second NLU result may be combined with the first NLU result generated in act 318 to produce a combined NLU result for interpretation by a speech-enabled application. In other embodiments, an NLU system may receive both the first ASR result and the second ASR result and the ASR results may be combined prior to performing NLU on the combined ASR result. In yet other embodiments, the second ASR result may be generated based, at least in part, on a portion of the first audio and at least a portion of the second audio, as described above. A benefit of these latter two approaches is that the ASR result processed by the NLU system appears to the NLU system as if it was recognized from a single utterance, and thus, may be more likely to generate an NLU result that can be interpreted by the speech-enabled application to perform a valid action.

[0044] As discussed above in connection with the process of FIG. 2, some embodiments are configured to process second audio in response to determining that received first audio does not include sufficient information to allow a speech-enabled

application to perform a valid action. Determining that the received first audio includes insufficient information for a speech-enabled application to perform a valid action may be made in any suitable way including, but not limited to, receiving an NLU result corresponding to the received first audio from an NLU system that includes such insufficient information.

[0045] The inventors have recognized and appreciated that storing or “caching” one or more prefixes in local storage accessible to a client device may facilitate the identification of audio that may include insufficient information to allow a speech-enabled application to perform a valid action if the default timeout value (e.g., three seconds) is used for endpointing. For example, commonly and/or frequently used prefixes known to be associated with user pauses may be locally stored by a client device, and identification of a locally-stored prefix in received audio may mitigate premature endpointing. Accordingly, in some embodiments, an ASR result corresponding to received audio may be compared to one or more prefixes stored locally on a client device, and a threshold time used for endpointing may be dynamically adjusted based, at least in part, on a threshold time associated with a matching locally-stored prefix, as discussed in more detail below. By dynamically adjusting the threshold time used for endpointing based on particular detected prefixes, the number of premature endpointing occurrences may be reduced. For example, prefixes that are often associated with user pauses following the prefix may be associated with longer threshold times, thereby allowing longer audio recording times for those utterances prior to endpointing. Additionally, because the threshold times are not lengthened for all utterances, the latencies associated with speech recognition and/or natural language processing are not substantially increased, resulting in a more desirable user experience than if the threshold times were increased for all utterances.

[0046] FIG. 4 illustrates a process for dynamically adjusting a threshold time used for endpointing in accordance with some embodiments. In act 410, audio is received (e.g., from a microphone). The process then proceeds to act 412, where the received audio is analyzed (e.g., by ASR processing and/or NLU processing) to determine whether the audio includes a locally-stored prefix. As discussed above, one or more prefixes may be locally stored on a client device configured to receive speech input. The locally-stored prefix(es) may include any suitable prefix often followed by a

pause that may cause an endpointing process to timeout. For example, certain prefixes such as “directions to” are often followed by pauses while the speaker thinks about what to say next. In some embodiments, the prefix “directions to” and/or other suitable prefixes may be stored locally on a client device.

[0047] The received audio may be processed and compared to the locally-stored prefix(es) in any suitable way, and embodiments are not limited in the respect. For example, at least a portion of received audio may be processed by an ASR engine, and an ASR result output from the ASR engine may be compared to locally-stored prefixes to determine whether the ASR result matches any of the locally-stored prefixes. The ASR processing and comparison should occur quickly enough to alter a timeout value for endpointing on the fly if a match to a locally-stored prefix is identified. For example, if the default timeout for endpointing is three seconds, the cached prefix lookup should preferably take less than three seconds to enable the timeout value used for endpointing to be lengthened, if appropriate, based on the identification of a cached prefix in the received audio.

[0048] In some embodiments, multiple short time segments (e.g., 20 ms) of the first audio may be processed by the ASR engine to facilitate the speed by which an ASR result is determined based, at least in part, on at least a portion of the received audio. As more of the audio is received, the ASR processing may continually update the ASR results output from the ASR engine. During this process, a current ASR result output from the ASR engine may be compared to the cached prefixes in an attempt to identify a match. Performing the ASR and comparison processes, at least in part, in parallel may further speed up the process of identifying a locally-stored prefix in the received audio.

[0049] If it is determined in act 412 that the received audio does not include a locally-stored prefix the process ends and the default timeout value for endpointing is used. If it is determined in act 412 that the received audio includes a locally-stored prefix, the process proceeds to act 414, where the timeout value for endpointing is dynamically set based, at least in part, on a threshold time associated with the identified locally-stored prefix.

[0050] Locally-stored prefixes and their associated threshold times may be stored in any suitable way using one or more data structures, and the one or more data

structures may be updated periodically and/or in response to a request to do so. The stored prefixes and their corresponding threshold times may be determined in any suitable way and may be user independent and/or user specific. For example, the locally-stored prefixes may be determined based, at least in part, on ASR data for a plurality of users to identify the most common prefixes that cause ASR to timeout. In some embodiments, an initial set of prefixes based on user-independent data analysis may be updated based on individual user behavior. In other embodiments, a user independent set of prefixes may not be used, and the local cache of prefixes may be determined by a manual selection or programming of the client device on which ASR is performed and/or the cache of prefixes may be established only after a user has used the client device for a particular amount of time to enable the cache to be populated with appropriate prefixes based on the user's individual behavior. In some embodiments, the locally-stored cache may include one or more first prefixes determined based on user-independent data analysis and one or more second prefixes determined based on user-specific data analysis. Not all embodiments require the use of both user-independent and user-specific prefixes, as some embodiments may include user-specific prefixes only, some embodiments may include user-independent prefixes only, and other embodiments may include both user-specific and user-independent prefixes.

[0051] As discussed above, each of the locally-stored prefixes may be associated with a threshold time suitable for use as an endpointing timeout value for the prefix. The threshold times for each prefix may be determined in any suitable way and one or more of the threshold times may be user independent and/or user specific, as discussed above. In some embodiments, a threshold time for a locally-stored prefix is determined based, at least in part, on an analysis of pause length for a plurality of speakers who uttered the prefix. For example, the threshold time for the prefix "directions to" may be determined based, at least in part, using the average pause length for 1000 utterances for different speakers speaking that prefix. The threshold time associated with a locally-stored prefix may be updated at any suitable interval as more data for a particular prefix is received from a plurality of speakers and/or from an individual speaker. By determining threshold times for prefixes from individual speakers, a suitable threshold time for each prefix may be established for the speaker, thereby providing an ASR system with reduced premature endpointing tuned to a particular speaker's speaking style.

[0052] Performing both ASR and NLU on a client device provides for low latencies and works even when a network connection to a server is not available. However, servers often include more processing and/or storage resources than client devices, which may result in better accuracy compared to client-based ASR processing and/or NLU processing. Hybrid ASR or NLU systems attempt to tradeoff accuracy for processing latency by distributing ASR and/or NLU processing between clients and servers in a client-server environment.

[0053] The inventors have recognized and appreciated that the increased latencies and intermittent server unavailability for server-based NLU systems are a contributing factor to user frustration with such systems. To improve user experiences with speech-based systems that perform at least some NLU processing on a server, some embodiments store locally on the client device, representations of recent and/or frequent utterances (e.g., ASR results for the utterances) and an NLU result associated with those representations. By locally caching NLU results for recent and/or frequent utterances, the corresponding NLU results may be quickly available on the client device even if the client device itself does not perform any NLU processing.

[0054] FIG. 5 illustrates a process for storing on a client device, in a client-server architecture, NLU results for one or more recent and/or frequent utterances in accordance with some embodiments. In act 510, an ASR result for first audio including speech is generated by an ASR process. The ASR process may be completely or partially performed using an ASR engine on the client device and/or the server. The process then proceeds to act 512, where an NLU process is performed by the server to generate an NLU result based, at least in part, on the ASR result. The NLU result is then returned to the client device.

[0055] After generating the NLU result, the process proceeds to act 514, where it is determined whether to store the generated NLU result in local storage associated with the client device. The determination of whether to locally cache the generated NLU result may be based on one or more factors including, but not limited to, how frequently the NLU result has been received from the server, the available storage resources of the client device, and how current the usage of the utterance is. For example, in some embodiments, provided that the client device has sufficient storage resources,

representations of all utterances and their corresponding NLU results for the previous 24 hour period may be cached locally on the client device. In other embodiments, representations of frequently recognized utterances within a particular period of time (e.g., five times in the past 24 hours) may be cached locally with their corresponding NLU results, whereas representations of less frequently recognized utterances within the same period of time (e.g., one time in the past 24 hours) and their NLU results may not be cached locally. Any suitable criterion or criteria may be used to establish the cutoff used in determining when to cache representations of utterances and their NLU results locally on a client device, and the foregoing non-limiting examples are provided merely for illustration.

[0056] If it is determined in act 514 not to locally cache the NLU result, the process ends. If it is determined in act 514 to locally cache the NLU result, the process proceeds to act 516, where a representation of the utterance (e.g., the ASR output associated with the utterance) and a corresponding NLU result return from the server is stored in local storage. For example, the representation of the utterance and its corresponding NLU result may be added to one or more data structures stored on local storage. In some embodiments, a small, local grammar that enables fast and highly-constrained ASR processing by the client device may be additionally be created for use in recognizing the frequently occurring utterance. It should be appreciated however, that not all embodiments require the creation of a grammar, and aspects of the invention are not limited in this respect.

[0057] FIG. 6 illustrates a process for using a locally-cached NLU result in accordance with some embodiments. In act 610, a client device performs an ASR process on audio including speech to produce an ASR result. The process then proceeds to act 612, where it is determined whether the ASR result includes any of the one or more representations of utterances locally stored by the client device.

[0058] If it is determined in act 612 that the ASR result does not include any of the cached utterance representations, the process ends. Otherwise, if it is determined in act 612 that the ASR result includes a locally-stored representation of an utterance, the process proceeds to act 614, where the cached NLU result associated with the identified locally-stored utterance representation is submitted to a speech-enabled application to

allow the speech-enabled application to perform one or more valid actions based, at least in part, on a locally-cached NLU result. For example, if the identified locally-stored utterance representation is “Call Bob at home,” the client device may access a contacts list on the client device to determine a home phone number for the contact “Bob,” and a phone call may be initiated by a phone application to that phone number. By caching at least some NLU results locally on a client device, the user experience with NLU-based applications on a client device is improved due to increased availability of the NLU results and reduced latencies associated with obtaining them. Additionally, one or more actions associated with some ASR results may be stored on a client device. By directly accessing the stored actions via the techniques described herein, a client device can appear to perform NLU processing for frequent utterances even if the client device does not have NLU processing capabilities.

[0059] Although the above-described examples of caching NLU results describe caching recent and/or frequent NLU results, it should be appreciated that any suitable type of NLU results may additionally or alternatively be cached locally on a client device. For example, in some embodiments, NLU results for utterances corresponding to emergency situations (e.g., “Call Fire Department”) may be locally cached due to their importance.

[0060] Some embodiments are directed to NLU-based systems that include speech-enabled applications, such as virtual assistant software, executing on a client device. The inventors have recognized and appreciated that users often have difficulty learning the available options for interacting with the some speech-enabled applications and often learn through trial and error or by reading release notes for the application. Some embodiments are directed to techniques for training users on the options for what they can say to a speech-enabled application while simultaneously helping users to complete an input in progress by providing real-time feedback by displaying dynamically-generated hints on a user interface of the client device.

[0061] FIG. 7 shows an illustrative process for displaying dynamically-generated hints on a user interface of a client device in accordance with some embodiments. In act 710, a first hint is created based, at least in part, on an ASR result determined for first audio. For example, if the ASR result is “make a meeting,” a first hint such as “make a

meeting on <day> from <start time> to <end time> titled <dictate> may be generated. The process then proceeds to act 712, where the first hint is displayed on a user interface of the client device. By displaying the first hint to the user of the client device, the client becomes aware of the structure and components of an utterance that the speech-enabled application is expecting to receive to be able to perform an action, such as scheduling a meeting.

[0062] After the first hint is displayed on the user interface, the process proceeds to act 714, where second audio comprising speech is received by the client device. For example, the user may say “make a meeting on Friday at one.” The process then proceeds to act 716, where a second hint is created based, at least in part, on an ASR result corresponding to the second audio. For example, continuing with this example, the second hint may be “make a meeting on Friday at 1 for <duration> titled <dictate>.” The process then proceeds to act 718, where the second hint is displayed on the user interface of the client device. By dynamically updating the hint displayed to the user, the user learns how to interact with the speech-enabled application and understands what additional information must be provided to the speech-enabled application in the current utterance to perform a particular action. Teaching users to say particular words, such as “titled” may facilitate the parsing of utterances to reliably separate the utterance into its component pieces, thereby improving an ASR and/or NLU process.

[0063] In some embodiments, at least some received second audio may be processed by an ASR engine and/or an NLU engine based, at least in part on a currently- or previously-displayed hint. For example, the second audio in the example above may be processed by an ASR engine using a grammar that restricts speech recognition to the components of the first hint, which may improve ASR accuracy for the second audio.

[0064] The above-described embodiments of the present invention can be implemented in any of numerous ways. For example, the embodiments may be implemented using hardware, software or a combination thereof. When implemented in software, the software code can be executed on any suitable processor or collection of processors, whether provided in a single computer or distributed among multiple computers. It should be appreciated that any component or collection of components that perform the functions described above can be generically considered as one or more

controllers that control the above-discussed functions. The one or more controllers can be implemented in numerous ways, such as with dedicated hardware, or with general purpose hardware (e.g., one or more processors) that is programmed using microcode or software to perform the functions recited above.

[0065] In this respect, it should be appreciated that one implementation of the embodiments of the present invention comprises at least one non-transitory computer-readable storage medium (e.g., a computer memory, a USB drive, a flash memory, a compact disk, etc.) encoded with a computer program (i.e., a plurality of instructions), which, when executed on a processor, performs the above-discussed functions of the embodiments of the present invention. The computer-readable storage medium can be transportable such that the program stored thereon can be loaded onto any computer resource to implement the aspects of the present invention discussed herein. In addition, it should be appreciated that the reference to a computer program which, when executed, performs the above-discussed functions, is not limited to an application program running on a host computer. Rather, the term computer program is used herein in a generic sense to reference any type of computer code (e.g., software or microcode) that can be employed to program a processor to implement the above-discussed aspects of the present invention.

[0066] Various aspects of the present invention may be used alone, in combination, or in a variety of arrangements not specifically discussed in the embodiments described in the foregoing and are therefore not limited in their application to the details and arrangement of components set forth in the foregoing description or illustrated in the drawings. For example, aspects described in one embodiment may be combined in any manner with aspects described in other embodiments.

[0067] Also, embodiments of the invention may be implemented as one or more methods, of which an example has been provided. The acts performed as part of the method(s) may be ordered in any suitable way. Accordingly, embodiments may be constructed in which acts are performed in an order different than illustrated, which may include performing some acts simultaneously, even though shown as sequential acts in illustrative embodiments.

[0068] Use of ordinal terms such as “first,” and “second” in the claims to modify a claim element does not by itself connote any priority, precedence, or order of one claim element over another or the temporal order in which acts of a method are performed. Such terms are used merely as labels to distinguish one claim element having a certain name from another element having a same name (but for use of the ordinal term).

[0069] The phraseology and terminology used herein is for the purpose of description and should not be regarded as limiting. The use of "including," "comprising," "having," “containing”, “involving”, and variations thereof, is meant to encompass the items listed thereafter and additional items.

[0070] Having described several embodiments of the invention in detail, various modifications and improvements will readily occur to those skilled in the art. Such modifications and improvements are intended to be within the spirit and scope of the invention. Accordingly, the foregoing description is by way of example only, and is not intended as limiting. The invention is limited only as defined by the following claims and the equivalents thereto.

CLAIMS

What is claimed is:

1. A computing device including a speech-enabled application installed thereon, the computing device comprising:
 - an input interface configured to receive first audio comprising speech from a user of the computing device;
 - an automatic speech recognition (ASR) engine configured to:
 - detect based, at least in part, on a threshold time for endpointing, an end of speech in the first audio; and
 - generate a first ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech; and
 - at least one processor programmed to:
 - determine whether a valid action can be performed by the speech-enabled application using the first ASR result; and
 - instruct the ASR engine to process second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the first ASR result.
2. The computing device of claim 1, wherein determining whether a valid action can be performed by the speech-enabled application using the first ASR result is based, at least in part, on a natural language understanding (NLU) result generated using the first ASR result.
3. The computing device of claim 2, wherein the processor is further programmed to submit the NLU result to the speech-enabled application, and wherein determining whether a valid action can be performed by the speech-enabled application using the first ASR result comprises receiving an indication from the speech-enabled application that a valid action cannot be performed in response to submitting the NLU result to the speech-enabled application.
4. The computing device of claim 1, wherein the input interface is further configured to receive the second audio, wherein the second audio includes audio

recorded after the detected end of speech in the first audio, and wherein the ASR engine is further configured to process the second audio.

5. The computing device of claim 4, wherein processing the second audio comprises:

determining whether the second audio includes speech; and

generating a second ASR result based, at least in part, on at least a portion of the second audio in response to determining that the second audio comprises speech.

6. The computing device of claim 5, wherein generating the second ASR result comprises generating the second ASR result based, at least in part, on at least a portion of the first audio and at least a portion of the second audio.

7. The computing device of claim 5, wherein the at least one processor is further programmed to:

determine whether a valid action can be performed by the speech-enabled application using a natural language understanding (NLU) result generated based, at least in part, on at least a portion of the first ASR result and at least a portion of the second ASR result; and

instruct the speech-enabled application to perform the valid action in response to determining that the valid action can be performed using the NLU result.

8. The computing device of claim 5, further comprising:

at least one storage device configured to store one or more prefixes, each of which is associated with a corresponding threshold time for endpointing; and

wherein determining whether a valid action can be performed by the speech-enabled application comprises determining whether the speech in the first audio includes a prefix of the one or more prefixes stored on the at least one storage device.

9. The computing device of claim 8, wherein the ASR engine is further configured to:

process a plurality of time segments of the first audio prior to detecting the end of speech in the first audio, and wherein determining whether the speech in the first audio includes a prefix stored on the at least one storage device comprises comparing output of the ASR engine determined based on the processed plurality of time segments to the one or more prefixes stored on the at least one storage device.

10. The computing device of claim 8, wherein the at least one processor is further programmed to:

update the threshold time used by the ASR engine for endpointing in response to determining that the speech in the first audio includes a prefix stored on the at least one storage device, wherein updating the threshold time comprises instructing the ASR engine to use the threshold time for endpointing associated with the prefix stored on the at least one storage device identified in the speech in the first audio to detect an end of speech in the first audio.

11. The computing device of claim 1, wherein the at least one processor is further programmed to:

create a first hint based, at least in part, on the first ASR result, wherein the first hint prompts the user for speech input corresponding to a valid action that can be performed by the speech-enabled application; and
present the first hint via a user interface of the computing device.

12. The computing device of claim 11, wherein the input interface is further configured to receive the second audio, wherein the ASR engine is further configured to process the second audio to generate a second ASR result, and wherein the at least one processor is further programmed to:

create a second hint based, at least in part, on the first ASR result and/or the second ASR result, wherein the second hint prompts the user for speech input corresponding to a valid action that can be performed by the speech-enabled application; and
present the second hint via a user interface of the computing device.

13. The computing device of claim 11, wherein presenting the first hint via the user interface comprises visually displaying the first hint on the user interface, and wherein the first hint hints of additional information to supplement the first audio to perform the valid action.

14. The computing device of claim 11, wherein the input interface is further configured to receive the second audio, wherein the ASR engine is further configured to process the second audio, wherein processing the second audio comprises performing ASR processing on the second audio based, at least in part, on information included in the first hint.

15. The computing device of claim 1, further comprising:

at least one storage device configured to store at least one data structure including information describing a plurality of natural language understanding (NLU) results and corresponding ASR output used to generate the plurality of NLU results;

wherein the at least one processor is further programmed to:

determine whether to add the first ASR result and a corresponding NLU result generated using the first ASR result to the at least one data structure stored on the at least one storage device; and

add the first ASR result and the corresponding NLU result generated using the first ASR result to the at least one data structure stored on the at least one storage device in response to determining that the first ASR result and the corresponding NLU result should be added.

16. The computing device of claim 15, wherein determining whether to add the first ASR result and the corresponding NLU result generated using the first ASR result to the at least one data structure comprises:

determining a number of times the corresponding NLU result has been received by the computing device from an NLU engine remotely located from the computing device; and

determining that the first ASR result and the corresponding NLU result should be added to the at least one data structure when the number of times the corresponding NLU result has been received by the computing device exceeds a threshold value.

17. The computing device of claim 15, wherein the input interface is further configured to receive third audio including speech from the user of the computing device, wherein the ASR engine is further configured to generate a second ASR result based, at least in part, on at least a portion of the third audio, and wherein the processor is further programmed to:

identify an ASR output stored in the at least one data structure corresponding to the second ASR result; and

submit the NLU result corresponding to the identified ASR output stored in the at least one data structure to the speech-enabled application to enable the speech-enabled application to perform an action based on the submitted NLU result.

18. The computing device of claim 17, wherein the at least one processor is programmed to submit the NLU result corresponding to the identified ASR output stored in the at least one data structure to the at least one data structure without sending a request for remote NLU processing of the third audio to an NLU engine remotely located from the computing device.

19. A method, comprising:

receiving, by an input interface of a computing device, first audio comprising speech from a user of the computing device;

detecting, by an automatic speech recognition (ASR) engine of the computing device, an end of speech in the first audio;

generating, by the ASR engine, an ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech;

determining whether a valid action can be performed by a speech-enabled application installed on the computing device using the ASR result; and

instructing the ASR engine to process second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the ASR result.

20. A computer-readable storage medium encoded with a plurality of instructions that, when executed by a computing device, performs a method, the method comprising:
- receiving first audio comprising speech from a user of the computing device;
 - detecting an end of speech in the first audio;
 - generating an ASR result based, at least in part, on a portion of the first audio prior to the detected end of speech;
 - determining whether a valid action can be performed by a speech-enabled application installed on the computing device using the ASR result; and
 - processing second audio when it is determined that a valid action cannot be performed by the speech-enabled application using the ASR result.




FIG. 1




FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/033736

A. CLASSIFICATION OF SUBJECT MATTER

INV. G10L15/08

ADD. G10L25/87 G10L15/04

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)

G10L

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

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

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 02/060162 A2 (ENTPR INTEGRATION GROUP INC [US]; BALENTINE BRUCE [US]; STRINGHAM REX) 1 August 2002 (2002-08-01) page 16, line 24 - page 19, line 31; figure 4 claim 8	1,19,20
A	US 2006/080096 A1 (THOMAS TREVOR [GB] ET AL) 13 April 2006 (2006-04-13) paragraph [0009] paragraphs [0080], [0081], [0085], [0087], [0089], [0091], [0094], [0102]; figures 4,8 ----- -/--	1,19,20



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

2 August 2016

Date of mailing of the international search report

11/08/2016

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Krembel, Luc

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2016/033736

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	DI LU ET AL: "Decision of response timing for incremental speech recognition with reinforcement learning", 2011 IEEE WORKSHOP ON AUTOMATIC SPEECH RECOGNITION AND UNDERSTANDING (ASRU 2011) : WAIKOLOA, HAWAII, USA, 11 - 15 DECEMBER 2011, IEEE, PISCATAWAY, NJ, 11 December 2011 (2011-12-11), pages 467-472, XP032126171, DOI: 10.1109/ASRU.2011.6163976 ISBN: 978-1-4673-0365-1 paragraphs [III.B], [III.Cn] -----	1,19,20
A	US 7 720 683 B1 (VERMEULEN PIETER J [US] ET AL) 18 May 2010 (2010-05-18) column 6, lines 4-18 column 14, lines 34-52 -----	1,19,20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/US2016/033736

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 02060162	A2	01-08-2002	AU 2002246550 A1 06-08-2002
			US 2004098253 A1 20-05-2004
			WO 02060162 A2 01-08-2002
US 2006080096	A1	13-04-2006	GB 2422279 A 19-07-2006
			US 2006080096 A1 13-04-2006
			US 2006080099 A1 13-04-2006
US 7720683	B1	18-05-2010	NONE