



(51) International Patent Classification:

G10L 15/22 (2006.01) *G10L 15/02* (2006.01)
G10L 15/20 (2006.01) *G06F 15/16* (2006.01)

(21) International Application Number:

PCT/US2016/033623

(22) International Filing Date:

20 May 2016 (20.05.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

14/752,450 26 June 2015 (26.06.2015) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).

(72) Inventors: **DEISHER, Michael**; 2665 NE Anna Avenue, Hillsboro, Oregon 97124 (US). **STEMMER, Georg**; Schleißheimer Strasse 59a, 80797 München (DE).

(74) Agents: **MALLIE, Michael J.** et al.; Blakely, Sokoloff, Taylor & Zafman LLP, 1279 Oakmead Parkway, Sunnyvale, California 94085 (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, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: LANGUAGE MODEL MODIFICATION FOR LOCAL SPEECH RECOGNITION SYSTEMS USING REMOTE SOURCES

FIG. 2



(57) Abstract: A language model is modified for a local speech recognition system using remote speech recognition sources. In one example, a speech utterance is received. The speech utterance is sent to at least one remote speech recognition system. Text results corresponding to the utterance are received from the remote speech recognition system. A local text result is generated using local vocabulary. The received text results and the generated text result are compared to determine words that are out of the local vocabulary and the local vocabulary is updated using the out of vocabulary words.



Published:

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

LANGUAGE MODEL MODIFICATION FOR LOCAL SPEECH RECOGNITION SYSTEMS USING REMOTE SOURCES

FIELD

The present description pertains to speech recognition for connected devices and in particular to speech recognition using both local and remote resources.

BACKGROUND

Cellular telephones and motor vehicles have long offered speech recognition systems for hands free operation, for navigation, and for controlling entertainment systems. These systems have suffered from the difficulty of understanding multiple languages, dialects, vocabularies, and pronunciation styles. Poor diction and background noise make speech recognition even more difficult. Some devices operate well by only recognizing a few statements. In some cases, the list of possible statements is displayed on a screen or recited audibly to the user by the system. The user makes one of the statements and then the device repeats the statement for confirmation.

With the advent of data connectivity for smart phones and the Internet of Things, large and powerful servers coupled to substantial databases are available to connected devices. This allows for much better recognition for more languages and more words. Newer products allow users to not only speak to their smart phone, television, and gaming console, but also to watches, fitness sensors, glasses and other portable and wearable devices.

With the increased use and variety of handheld and wearable communications devices, speech understanding and audio quality have become increasingly important. Many handheld and wearable devices receive speech and other audio and send the captured audio to a remote server. The remote server converts the speech to text or commands and sends it back to the connected device. This allows the speech to be used for voice commands. Cloud speech recognition systems are designed to provide good accuracy independent of the acoustics, vocabulary, and grammar of the user. Systems implemented in the cloud can use very large acoustic and language models to meet this goal and can update models frequently.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments are illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings in which like reference numerals refer to similar elements.

Figure 1 is a block diagram of a system of connected devices with access to remote resources according to an embodiment.

Figure 2 is a block diagram of operation of the connected devices of Figure 1 according to an embodiment.

Figure 3 is a process flow diagram of recognizing speech and updating a local speech recognition system according to an embodiment.

Figure 4 is a process flow diagram of identifying out-of-vocabulary words at a local speech recognition system according to an embodiment.

Figure 5 is a process flow diagram of an alternative for identifying out-of-vocabulary words at a local speech recognition system according to an embodiment.

5 Figure 6 is a block diagram of a computing device incorporating speech recognition and local speech recognition enhancement according to an embodiment.

DETAILED DESCRIPTION

Notwithstanding the availability of remote connections, server-based automated speech recognition systems have at least four shortcomings. First, the network may not always be
10 available. Second, a network radio such as Wi-Fi or cellular uses power, reducing battery life for a portable device. Third, some users may want to avoid sending text, commands, calendar, and contact information to a server for privacy reasons. Fourth, interposing two-way communication and server processing may slow down the user interface. However, when the device is not using the powerful server system, the speech recognition user experience is degraded. This is because
15 of the limited local resources on the client device. This includes limited client processing capability and memory limitations.

Cloud speech recognition systems can use very large acoustic and language models update models frequently. Large server systems with a large database of clients can also use predictive systems to start processing a speech utterance before it is completed. Very common
20 questions can be answered very quickly. Even if uncommon questions take longer, the user experience is mostly improved. For a client system, it is usually sufficient to provide good accuracy to a single user speaking a single language and dialect and having a similar style of pronunciation and diction. This reduced variety allows a compact client model to be constructed with a smaller amount of memory that performs as well or better for a single or a few users than
25 does a large user-independent model in the cloud. Even when a device has multiple users, since they are often related by geographic location, workplace, or family, the client model may still remain very small.

One or more ASR (Automated Speech Recognition) systems in the “cloud” may be used to improve the accuracy of a client speech recognition system. Cloud speech recognition may be
30 done parallel with client speech recognition and results may be compared. Words that are Out-Of-Vocabulary (OOV) for the client device may be identified by the cloud ASR and added to the client lexicon. This may be done using a dynamic vocabulary mechanism or using opportunistic model rebuilding.

The client system may be further enhanced by confidence mechanisms. User behavior in
35 response to cloud recognition results may be used to estimate a confidence for the accuracy of a

particular cloud speech recognition system. The client device's local language model n-gram frequencies or neural network weights may be updated to reflect high confidence user utterances. In addition, the language model may be pruned and grown based on these updated weights. This allows the local memory to be increasingly better tailored to the actual users of the device. The language model may then stay within the client device's memory size constraints. In addition, the local processors are able to provide faster results using a more compact and relevant language model.

Figure 1 is a diagram of a connected device system in which a variety of different user devices communicate through the cloud to various servers and service providers. On the client device side, there is a wearable device such as glasses or a headset 102, a watch 104, a wearable fitness sensor in a shoe, strap, or other article of clothing 106, a computing tablet 108, and a portable handset 110 such as a smart phone, navigation device, fitness device, etc. The client side may also include media players, televisions, set-top boxes, gaming devices, computers, workstations, and other devices. These devices all may communicate with each other as suggested by the arrows using a local networking system, such as Wi-Fi, Bluetooth, or another wired or wireless protocol. The devices also communicate either directly or through another one of the devices through a wider area network 112, such as a WAN (Wide Area Network), MAN (Metropolitan Area Network), LAN (Local Area Network), the Internet or any other suitable communications system.

On the service side there may be one or more Automated Speech Recognition systems 122, 124, 126, or other similar voice recognition and command systems. There may also be numerous service providers 128. These service providers may serve by purchasing tickets or other items, updating cloud-based or local device scheduling, calendaring or other storage systems, sending communications to other users, and in other ways.

Figure 2 is a diagram of the operation of any one or more of the client devices of Figure 1 in cooperation with one or more ASRs or similar services. In the example of Figure 2, the client device has a microphone 202. The microphone receives an utterance from a user (not shown). Typically, the microphone converts the utterance to a digital form for transmission to other devices and through the network, however, this is not required. There is a direct path 204 from the microphone to a client application 206. This path may be within a single client device, for example from a built-in microphone on a headset or a watch to a processor on the same device, or it may be from one client device to another. The client application 206 accesses a client device automated speech recognition (ASR) system 208 to analyze the utterance. The ASR may be part of the client app 206 or it may be separate module accessed using an application programming interface, operating system calls or in any other desired way.

The speech utterance is also conveyed from the microphone through a network path 222 to one or more remote short dictation applications 224 coupled to corresponding remote ASRs 226, such as a cloud ASR. The network path may be through the same or another client device. As an example, a watch may send voice commands to a smart phone which sends the captured
5 utterance to various servers through a cellular or Wi-Fi link to the internet. A headset may do the same or send the commands to a fixed gaming system which is connected to a wired internet connection to remote servers. A variety of different paths may be used. In some cases, wearable and other devices have a built-in connection to remote networks or servers and these may be used directly.

10 The client ASR 208 may be connected to and access a variety of different resources such as an acoustic model 210, and a language model 212. The language model may be connected to and access a database of N-gram counts 214, including bigrams, trigrams, and higher count grams, and a lexicon 216. The language model may have many additional resources, depending on the particular implementation. The client device ASR uses these tools to generate a text result
15 corresponding to the utterance received at the microphone. This utterance may then be used by the client app as commands or information. The commands may be for storing data, adding or changing records in contact, calendars, notes or other systems. The commands may be for sending messages to other users or to perform transactions on other systems. The commands may also be to operate components of the client device or another connected device.

20 The cloud ASR may also be connected to and access an acoustic model 228 and a language model 230 to generate a text result for the utterance. In general the acoustic model and language model is significantly larger and able to interpret significantly more different utterances from more varied users. The generated text result is sent back to the client device through the same or another network path 234. As shown this result is received directly at the client device
25 ASR for use in making a determination of the final text result using both the locally generated text and the remotely generated text result. The result may be received by the same device that houses the ASR, by another connected device or in any other way. The text result will be transferred to the client device ASR through a software stack that recognizes the active client application 206. The client application will receive the final selected text result from the local
30 ASR for the utterance.

The client ASR compares the received text result and the generated text result to determine a final result and then provides the final result to the client application. The client application then performs an action based on the final text result. This may be an action performed by the client device or the action may be to send the final text result to another
35 application, module, or device. While only one remote ASR is shown for simplicity, more than

one remote ASR may be used and the client ASR may combine these to obtain better, faster, or more accurate results. The client ASR may also compare results to determine which remote ASRs provide the best results and then select those in favor of other ASRs that provide worse results.

5 Using several cloud ASR systems will increase local power consumption and could increase the delay experienced by the end user. This will make the device respond more slowly to any spoken utterance. In some embodiments, the local device uses the results from a cloud ASR directly without any comparison or local ASR. This may provide higher accuracy at the cost of some delay for the remote communication. In some embodiments, the device uses only
10 the locally generated results. This provides a quicker response. In such cases, the cloud ASR results may be used for improving the client ASR instead of for responding to the user speech utterance. The process of improving the client ASR, as described herein may use additional power beyond what is required to respond to the utterance. As a result, this process may be turned off when the device is in a power saving mode, whether to conserve power or when
15 battery levels are low. As a further power-saving measure the remote cloud access may be turned off, even when a remote connection is available.

Figure 3 is a process flow diagram of recognizing a speech utterance and updating a local speech recognition system according to an embodiment. Such a process may be implemented using the hardware of Figure 2 and variations thereof, *inter alia*. At 302, a speech utterance is
20 received. This is typically in the form of digitized audio from a local microphone. However, the speech utterance may be received in other ways. The speech utterance may be packetized and provided to a buffer by an audio processing pipeline. The audio of the utterance may be recorded into a buffer or longer term memory at the local device for further analysis. At the buffer, it is available to a main processor or to a separate speech recognition system. In order to
25 interpret and speech utterance, the system will use one or more automated speech recognition systems.

At 308, the audio of the speech utterance is sent to the local ASR. This operation may be performed every time for any speech utterance in order to use local results. However, if cloud ASRs are available, then they may be used instead and this operation may be skipped.

30 At 304, the system determines whether it is connected to a network. The system may not be connected because there is no network available or because the network interface is not available. The network interface may be turned off to save power or to prevent interference caused by transmissions such as in an "airplane" mode or for any other reason. If there is no network connection, then the system performs a local ASR and at 330 waits for the local ASR to

finish. After the local ASR is finished, then this result is applied locally to the appropriate application at 332. The system is operating autonomously using the available local ASR.

If there is a network connection, then the system still performs a local ASR at 330 as it continues to evaluate connections and remote ASR operation. As an alternative, the system may be configured to use only remote ASRs if remote ASRs are available. As another alternative, the system always performs a local ASR at 330 and uses the result for the application at 332. Cloud results, if any, are used only for updating the local ASR at 326.

Similarly at 306, the system determines whether there are any user settings to restrict network connections. The system may be set not to use a network connection even when such a connection is available. This may be for privacy purposes, to reduce data transmission, to conserve battery power, or for any other reasons. In some embodiments, the system is configured to primarily rely on the local ASR in the user setting restrictions. The cloud ASRs are used only in certain specified circumstances. In one example, the system may be set to use only Wi-Fi connections and not to use cellular telephone connections in order to reduce data costs. Alternatively, the system may be set to not use Wi-Fi connections for security or privacy purposes. If there are restrictions on the use of the data network, then the system performs a local ASR at 330 and uses the results for the active application at 332. The local ASR converts the speech utterance to text and this text is used by the active application as data, command, or in any other way, depending on the nature or the active application.

If a network connection is available and there are no restrictions on use of the network, then at 310, the system sends the audio to one or more network connected ASR's. These systems will be referred to as the cloud ASR's, however, they may be accessible through any of a variety of different types of network connection and may be internal or external resources for a network or other system. The cloud ASRs, like the local ASR, analyze the speech utterance as one or more digitized audio samples and then generate some type of machine recognizable data. Typically the data is a text sequence but could be in any of a variety of other forms.

The remote speech recognition systems may be selected based on pre-determined parameters, user selection or based on past performance. The ASR systems that provide the best results may be returned to for later utterances. In this way the local device is optimized for the particular users and for the available resources which may change over time.

The cloud ASR interaction may be managed using different tests. First at 312, there is a timer. When the timer times out then the system determines whether all of the cloud ASR results are available at 316. If the cloud ASR results are all available then the system determines whether the cloud ASR results all match at 318. If the cloud results all match then the system proceeds at 330 to wait for the local ASR results. After these are received, then at 332 the local

ASR-generated text is used by the application and the speech recognition process ends. If the results do not match, then the system retrieves the local ASR results at 334 and then selects the best text at 320.

If the timer has not timed out, then the system may still determine whether all of the cloud ASR results have been received at 316. If they have all been received, then the system may proceed to determine whether they all match at 318. If the results have not all been received at 316 and the timer has not expired at 312, then the system continues to wait until the timer expires before operating on the cloud ASR results.

After the timer expires at 312 or if all of the results have been received at 316, then all of the cloud ASR results that have been received may be compared at 318. The text results received from the remote speech recognition system correspond to the original user speech utterance. While the selected remote ASRs are working, the local ASR may also generate a text result corresponding to the utterance at 334. The local ASR then compares the received text results and the generated text result. These results may be combined based on the comparison to determine a final text result for the utterance at 320. The final result may then be used by a client application at 322 as a command, as data, or in a variety of other ways.

In addition the final result may be used to update the local acoustic model and the local language model at 326. When the received remote ASR results do not match, then at 320 the best text string is selected or created using acoustic confidence weighting or in some other way.

In some circumstances, whether the cloud ASR results match or not, the remotely received text result may or may not be used by the client application if the local speech recognition result confidence is high. In addition, there will be no remote text results unless the privacy permissions at the client device allow the speech utterance to be sent in each particular context. If the local and remote cloud speech recognition results differ, then the differences are analyzed. This may be done in a local phone recognition lattice or any other analysis system available to the local ASR.

At 320, if the cloud ASR results are to be used by the application, then the best text string between the cloud ASR and the local ASR results is selected. This may be done with acoustic confidence weighting or any other desired type of scoring. In some cases, in order to provide a faster response, the local ASR result is used regardless of the cloud ASR results. At 322 once a text result is selected, then this result is applied to the active application.

At 324, given the local ASR and the cloud ASR results, out of vocabulary (OOV) words are identified. That is the system examines the local and cloud results to determine if there are any words in the cloud ASR results which are not in the local ASR. At 326, if any OOV words are found, then when appropriate the local ASR is updated at 326 using the identified OOV

words. The system then ends the process at 328 until another speech utterance is received at 302.

Figure 4 is a process flow diagram of how OOV words may be identified according to some embodiments. Figure 4 shows an example of how the operations at 324 may be performed.

5 By adding OOV words to the local ASR, the quality of the results produced by the local ASR may be improved using the remote resources. The language model, acoustic model, n-gram counts, and the lexicon, as well as any other appropriate resources of the local ASR may be augmented by the cloud ASR results. The improvement of the local ASR may be done after the ASR results are provided to the active application in a background process. The improvement
10 may also be done only when there is sufficient battery power. This improvement process is only done when remote ASR results are available. As mentioned above, the remote ASR results may or may not be used in the speech recognition process. The system may use local ASR for speech recognition and then use the remote results only to improve the local ASR by adding OOV words and any other desired data.

15 Different processes may be used, depending on the circumstances. If there is only one cloud ASR system available for results. Then the results from that one cloud system may be assumed to be the single-best hypothesis. Depending on settings and delay, this hypothesis may also be applied to the active application. If there is more than one cloud ASR result available, then voting or other methods may be used to either choose one of the results or to construct one
20 single-best hypothesis from the received results. Multiple candidate text strings may be constructed using combinations of the results to then select the best single text string.

Referring to Figure 4, at 402, the received speech utterance is evaluated to determine an acoustic confidence level. This test may be used as a gate test to determine whether the utterance should be used for enhancing the local ASR. Any of a variety of different acoustic
25 confidence scoring techniques may be used such as normalized acoustic likelihood, posterior probability, spread of n-best scores, etc. If the acoustic confidence is too low, then there may be an error in the received utterance due to noise, interference, or some other muffling or blocking effect. In some cases, the received acoustics may be clearly received, however, there is no suitable match in the language model. Such an acoustic confidence score may also be used in an
30 ASR system to reject a user speech utterance and request that the user repeat the utterance before any further voice recognition processing is performed. In that case, the confidence required to require the user to repeat the utterance may be much lower than the confidence required to use the utterance to determine OOV words.

If the acoustic confidence is too low for the local ASR enhancement then the process
35 ends at 418. If the acoustic confidence is sufficiently high then at 404 selection criteria are

applied to all of the cloud results to select the best one. If there is only one cloud ASR result, then this is used and selected as the best. With multiple received text strings from the cloud ASR's, the evaluations may include voting, neural networks and other evaluation systems.

The determination of high or low may be done using one or more thresholds. There may
5 be a low threshold so that if the result is below the threshold then the results are rejected. There may be a high threshold so that if the results are above or exceed the high threshold, then the results are used to possibly update the local ASR. If the acoustic confidence is between the thresholds, then the system may return a speech recognition result but not modify the local ASR. Different threshold configuration and application systems may be used depending on the
10 particular implementation.

At 406, the selected cloud text string is compared to the local ASR text string to determine whether there are any words in the cloud text string that are not in the local text string. In some audio speech samples, it is not always clear where one word ends and the next begins. For this and other reasons, differences between the cloud and local text strings may not be due to
15 a difference in lexicons, but due to a difference in interpreting the utterance as words. String matching algorithms and other techniques may be used to identify differences in the words between the two text strings. This may be done using standard insertion or substitution identification algorithm, such as NIST (National Institute of Standards and Technology) sctk (scoring toolkit) or a variety of other techniques. A proper substitution from the cloud string into
20 the local string may be identified as an OOV word in the local ASR.

In some embodiments, the identified OOV words cannot be directly added into the lexicon because their true pronunciation is not known. At 408 pronunciations are generated for the identified OOV words. This may be done using a grapheme to phoneme (G2P) algorithm or the pronunciations may be looked up in a lexicon. These pronunciations may be inserted into a
25 local phone lattice. At 412, the revised local phone lattice may then be scored again using the substituted words. This is another acoustic scoring in which the local phone lattice built from the text string is compared to the audio received from the user.

If the acoustic score for the new text string is low, then the results are rejected. In some embodiments, the process ends as shown. In other embodiments, the local phone recognition
30 lattice is used to resolve pronunciation ambiguities. Different words from the remote text string may be substituted into the local text string to generate different candidate text strings. These possibilities for the local text string are scored and the one with the highest score is selected. If the acoustic confidence is still too low, then the process ends. The local ASR is not changed.

On the other hand, if the acoustic confidence score is high, either with a single text string
35 or with one selected from multiple candidates, then the local ASR may be improved. Different

aspects of the local ASR may be modified, depending on the nature of the local ASR and the particular implementation. In some embodiments, as shown at 416, the n-gram counts and the local lexicon is enhanced by adding the newly identified OOV words from the new text string. The n-gram counts may be used for future language model updates, among other purposes.

5 In some embodiments, the update of local ASR may be deferred. The update may be done at the client when resources allow, or the update may be done at a cloud service. A cloud service update may be desired, for example, when the local device is simply too small or too low power to perform the update or to provide central control. In one embodiment, in order to perform the update of the local language model, in the cloud, the OOV list, n-gram counts, any
10 other desired data may be sent to a cloud service. The cloud service may then generate the new models and then transmit them to the client when the network allows. In this way the local device is only writing the new data to its memory instead of making the determinations of proper values for the new data.

 Figure 5 is an alternative process for identifying OOV words and then augmenting the
15 local ASR using the identified words. In the example of Figure 5, more than one remote or cloud ASR has provided a text result and the result are, at least in some cases, different from each other. As in the previous example, the acoustic confidence of the original speech utterance may be assessed at 502. If the confidence is too low, then the process ends at 518. If the confidence is sufficiently high, then at 504, a string matching algorithm is used to identify words
20 in each text string that may be OOV words. As mentioned above, high or low may be evaluated using numerical thresholds or in any other desired way. There may be one or multiple thresholds, depending on the particular implementation.

 Because there are multiple different text strings, there are more candidate words to consider. String matching and other types of techniques may be used to identify the differences
25 between the strings. The inserted or substituted words in each cloud text string may then be identified as possible OOV words.

 At 506 pronunciations may be generated for these identified OOV words. As mentioned above, this may be done using G2P or a lexicon or in any other desired way. At 508, each different text string or hypothesis from the cloud ASRs may be tested by lattice rescoring or in
30 any of a variety of other ways. Substitution or insertion of OOV words from the various cloud text strings into the local text string may be used to form multiple candidate text strings. In some embodiments, this may be done with a phone lattice that includes all of the possible variations. The phone lattice may be used by substituting or inserting OOV word pronunciations into the local phone lattice to modify the local phone lattice to generate multiple cloud hypotheses.

At 510, the newly modified local phone lattice is rescored for each variation within the lattice. Alternatively, various candidates may be generated by substitution and each scored using an acoustic confidence score. Alternatively, based on this scoring, a text string can be selected as the best candidate. This may be a particular text string received from the cloud or it may be a combination of different text strings received from different cloud ASRs.

At 512 the text string hypothesis with the highest score is selected. Using the local phone lattice, each hypothesis of the modified local phone lattice will be tested against the actual utterance by scoring and the hypothesis with the highest score will be selected. In one embodiment, the cloud ASRs are only used to substitute low scoring words in the local ASR result with words from the cloud ASRs. In another embodiment, only the cloud ASRs are used to form the hypotheses. A phone lattice allows many different hypotheses to be tested through the single lattice structure.

At 514, the acoustic confidence for the highest scoring candidate text string is evaluated. If the confidence is too low, then the process ends and the local ASR is unchanged. If the confidence is sufficiently high, then at 516 the local ASR may be updated for the OOV words as described above. This may include updating the n-gram counts for a later language model update, updating the OOV list for a later lexicon update, and updating other features.

When updating the local ASR at 416 or 516, a variety of different techniques may be used. The final result in either process results in one or more OOV words being identified. These final results may be used to update the local acoustic model and the local language model, *inter alia*. As mentioned above with respect to 416, the updating may be deferred based on power or other system usage. The updating may also be performed using a cloud service to actively determine the updates.

In some embodiments, the local acoustic model 210 and the local language model 212 are pre-installed for each market and are never changed except with other general updates. In some embodiments, the acoustic model and the language model including the lexicon are updated using the received text results from the remote ASRs. In this way, the original data is updated based on new data and results.

The client lexicon and the language model may be augmented using a cached LM technique or by interpolation. The lexicon may be updated directly, for example by rebuilding a HCL (Hidden Markov model Context model Lexicon grammar) in a WFST-based (Weighted Finite State Transducer) system. In some embodiments, the lexicon is updated indirectly using dynamic vocabulary mechanisms.

In addition to updating the lexicon, the language model may also be rebuilt using user-specific client data (e.g. n-gram counts) from the final text results. This may be done at other

times, when the ASR resources are not being used for speech recognition in order not to interfere with perceived speed. In addition, after the language model has been updated, the acoustic model may also be rebuilt when resources are available using alignments produced with the updated language model.

5 The local lexicon is updated with these OOV words that are remotely received. The n-gram frequencies are also updated. After some use, if there are some words that are not used or are used only infrequently by the device users, then these may also be removed from the local lexicon and the language model.

Figure 6 is a block diagram of a computing device 100 in accordance with one
10 implementation. The computing device 100 houses a system board 2. The board 2 may include a number of components, including but not limited to a processor 4 and at least one communication package 6. The communication package is coupled to one or more antennas 16. The processor 4 is physically and electrically coupled to the board 2.

Depending on its applications, computing device 100 may include other components that
15 may or may not be physically and electrically coupled to the board 2. These other components include, but are not limited to, volatile memory (e.g., DRAM) 8, non-volatile memory (e.g., ROM) 9, flash memory (not shown), a graphics processor 12, a digital signal processor (not shown), a crypto processor (not shown), a chipset 14, an antenna 16, a display 18 such as a touchscreen display, a touchscreen controller 20, a battery 22, an audio codec (not shown), a
20 video codec (not shown), a power amplifier 24, a global positioning system (GPS) device 26, a compass 28, an accelerometer (not shown), a gyroscope (not shown), a speaker 30, a camera 32, a microphone array 34, and a mass storage device (such as hard disk drive) 10, compact disk (CD) (not shown), digital versatile disk (DVD) (not shown), and so forth). These components may be connected to the system board 2, mounted to the system board, or combined with any of
25 the other components.

The communication package 6 enables wireless and/or wired communications for the transfer of data to and from the computing device 100. The term “wireless” and its derivatives may be used to describe circuits, devices, systems, methods, techniques, communications channels, etc., that may communicate data through the use of modulated electromagnetic
30 radiation through a non-solid medium. The term does not imply that the associated devices do not contain any wires, although in some embodiments they might not. The communication package 6 may implement any of a number of wireless or wired standards or protocols, including but not limited to Wi-Fi (IEEE 802.11 family), WiMAX (IEEE 802.16 family), IEEE 802.20, long term evolution (LTE), Ev-DO, HSPA+, HSDPA+, HSUPA+, EDGE, GSM, GPRS, CDMA,
35 TDMA, DECT, Bluetooth, Ethernet derivatives thereof, as well as any other wireless and wired

protocols that are designated as 3G, 4G, 5G, and beyond. The computing device 100 may include a plurality of communication packages 6. For instance, a first communication package 6 may be dedicated to shorter range wireless communications such as Wi-Fi and Bluetooth and a second communication package 6 may be dedicated to longer range wireless communications such as GPS, EDGE, GPRS, CDMA, WiMAX, LTE, Ev-DO, and others.

The microphones 34 and the speaker 30 are coupled to an audio front end 36 to perform digital conversion, coding and decoding, and noise reduction as described herein. The processor 4 is coupled to the audio front end to drive the process with interrupts, set parameters, and control operations of the audio front end. Speech recognition processing may be performed in the communication package 6, the processor 4 or in any other device.

In various implementations, the computing device 100 may be eyewear, a laptop, a netbook, a notebook, an ultrabook, a smartphone, a tablet, a personal digital assistant (PDA), an ultra mobile PC, a mobile phone, a desktop computer, a server, a set-top box, an entertainment control unit, a digital camera, a portable music player, or a digital video recorder. The computing device may be fixed, portable, or wearable. In further implementations, the computing device 100 may be any other electronic device that processes data.

Embodiments may be implemented as a part of one or more memory chips, controllers, CPUs (Central Processing Unit), microchips or integrated circuits interconnected using a motherboard, an application specific integrated circuit (ASIC), and/or a field programmable gate array (FPGA).

References to “one embodiment”, “an embodiment”, “example embodiment”, “various embodiments”, etc., indicate that the embodiment(s) so described may include particular features, structures, or characteristics, but not every embodiment necessarily includes the particular features, structures, or characteristics. Further, some embodiments may have some, all, or none of the features described for other embodiments.

In the following description and claims, the term “coupled” along with its derivatives, may be used. “Coupled” is used to indicate that two or more elements co-operate or interact with each other, but they may or may not have intervening physical or electrical components between them.

As used in the claims, unless otherwise specified, the use of the ordinal adjectives “first”, “second”, “third”, etc., to describe a common element, merely indicate that different instances of like elements are being referred to, and are not intended to imply that the elements so described must be in a given sequence, either temporally, spatially, in ranking, or in any other manner.

The drawings and the forgoing description give examples of embodiments. Those skilled in the art will appreciate that one or more of the described elements may well be combined into a

single functional element. Alternatively, certain elements may be split into multiple functional elements. Elements from one embodiment may be added to another embodiment. For example, orders of processes described herein may be changed and are not limited to the manner described herein. Moreover, the actions of any flow diagram need not be implemented in the order shown; nor do all of the acts necessarily need to be performed. Also, those acts that are not dependent on other acts may be performed in parallel with the other acts. The scope of embodiments is by no means limited by these specific examples. Numerous variations, whether explicitly given in the specification or not, such as differences in structure, dimension, and use of material, are possible. The scope of embodiments is at least as broad as given by the following claims.

The following examples pertain to further embodiments. The various features of the different embodiments may be variously combined with some features included and others excluded to suit a variety of different applications. Some embodiments pertain to a method that includes receiving a speech utterance at a local device, sending the speech utterance to at least one remote speech recognition system, receiving a text result corresponding to the utterance from the remote speech recognition system, generating a text result corresponding to the utterance at the local device using a local language model including a local vocabulary in a local lexicon, comparing the received text result and the generated text result to determine words that are out of the local vocabulary, and updating the local language model using the out of vocabulary words.

Further embodiments include combining the generated and received text results based on the comparison to determine a final text result for the utterance and determining the out of vocabulary words using the final text result.

In further embodiments combining comprises using a phone lattice on the local device if the generated and the received text results differ.

In further embodiments comparing comprises applying a neural network.

Further embodiments include determining an acoustic confidence value for the generated text result at the local device and applying voting to the generated and received text results if the acoustic confidence value exceeds a threshold to determine a text candidate for each word.

In further embodiments combining the generated and received text results comprises selecting a word from the client device lexicon unless a respective word is not available in the client device lexicon.

In further embodiments comparing comprises scoring words in the generated and received text result and substituting words from the received text result into the generated text result based on the scores.

In further embodiments comparing comprises inserting different words from multiple received remote cloud text strings into a phone lattice and scoring the phone lattice for each combination.

In further embodiments updating the language model comprises updating the client
5 device lexicon using the received text results.

In further embodiments updating the local language model comprises updating using out of vocabulary words from the final text utterance.

In further embodiments updating the local language model further comprises updating n-gram counts of the local language model.

10 Further embodiments include determining a speech recognition confidence value for the received text result, comparing the speech recognition confidence value to a threshold and not updating the local vocabulary if the speech recognition confidence value for the received text result does not exceed the threshold.

Further embodiments include determining at the client device whether the device is
15 connected to a remote network and wherein sending the audio utterance is performed only if the device is connected to a remote network.

Further embodiments include conditionally sending the speech utterance to a remote speech recognition system based on user privacy permissions for the local device.

Some embodiments pertain to an apparatus that include a buffer of a local device to
20 receive and record a speech utterance, a network interface to send the speech utterance to at least one remote speech recognition system and to receive a text result corresponding to the utterance from the remote speech recognition system, a processor to generate a text result corresponding to the utterance at the local device using a local language model including a local vocabulary in a local lexicon, to compare the received text results and the generated text result to determine
25 words that are out of the local vocabulary, and to update the local language model using the out of vocabulary words.

In further embodiments updating the local language model comprises updating the local lexicon and n-gram counts of the local language model using the received text results.

In further embodiments the processor is further to determine an acoustic confidence value
30 for the generated text result at the local device, and to apply voting to the generated and received text result if the acoustic confidence value exceeds a threshold to determine a text candidate for each word.

Some embodiments pertain to a computing system that includes a microphone, a buffer of a local device to receive and record a speech utterance from the microphone, a local speech
35 recognition system to convert the speech utterance to a text result using a local language model

including a local vocabulary in a local lexicon, a client application to receive the text result as a command, a network interface to send the speech utterance to at least one remote speech recognition system and to receive a text result corresponding to the utterance from the remote speech recognition system, a processor to compare the received text results and the local text result to determine words that are out of the local vocabulary, and to update the local language model using the out of vocabulary words.

Further embodiments include combining the generated and received text results based on the comparison to determine a final text result for the utterance and determining the out of vocabulary words using the final text result.

In further embodiments comparing comprises inserting different words from multiple received remote cloud text strings into a phone lattice and scoring the phone lattice for each combination, and substituting words from the received text result into the generated text result based on the scores.

What is claimed is:

1. A method comprising:
receiving a speech utterance at a local device;
sending the speech utterance to at least one remote speech recognition system;
5 receiving a text result corresponding to the utterance from the remote speech recognition system;
generating a text result corresponding to the utterance at the local device using a local language model including a local vocabulary in a local lexicon;
comparing the received text result and the generated text result to determine words that
10 are out of the local vocabulary; and
updating the local language model using the out of vocabulary words.

2. The method of Claim 1, further comprising combining the generated and received text results based on the comparison to determine a final text result for the utterance and determining the out of vocabulary words using the final text result.

15 3. The method of Claim 2, wherein combining comprises using a phone lattice on the local device if the generated and the received text results differ.

4. The method of Claim 3, wherein comparing comprises applying a neural network.

5. The method of any of claims 2-4, further comprising:
determining an acoustic confidence value for the generated text result at the local device;
20 and
applying voting to the generated and received text results if the acoustic confidence value exceeds a threshold to determine a text candidate for each word.

6. The method of Claim 5, wherein combining the generated and received text results comprises selecting a word from the client device lexicon unless a respective word is not
25 available in the client device lexicon.

7. The method of any of claims 2-6, wherein comparing comprises scoring words in the generated and received text result and substituting words from the received text result into the generated text result based on the scores.

8. The method of Claim 3, wherein comparing comprises inserting different words
30 from multiple received remote cloud text strings into a phone lattice and scoring the phone lattice for each combination.

9. The method of any one or more of the above claims, wherein updating the language model comprises updating the client device lexicon using the received text results.

10. The method of Claim 7, wherein updating the local language model comprises
35 updating using out of vocabulary words from the final text utterance.

11. The method of any one or more of the above claims, wherein updating the local language model further comprises updating n-gram counts of the local language model.

12. The method of any one or more of the above claims, further comprising determining a speech recognition confidence value for the received text result, comparing the speech recognition confidence value to a threshold and not updating the local vocabulary if the speech recognition confidence value for the received text result does not exceed the threshold.

13. The method of any one or more of the above claims, further comprising determining at the client device whether the device is connected to a remote network and wherein sending the audio utterance is performed only if the device is connected to a remote network.

14. The method of any one or more of the above claims, further comprising conditionally sending the speech utterance to a remote speech recognition system based on user privacy permissions for the local device.

15. An apparatus comprising:

a buffer of a local device to receive and record a speech utterance;

a network interface to send the speech utterance to at least one remote speech recognition system and to receive a text result corresponding to the utterance from the remote speech recognition system; and

a processor to generate a text result corresponding to the utterance at the local device using a local language model including a local vocabulary in a local lexicon, to compare the received text results and the generated text result to determine words that are out of the local vocabulary, and to update the local language model using the out of vocabulary words.

16. The apparatus of Claim 15, wherein updating the local language model comprises updating the local lexicon and n-gram counts of the local language model using the received text results.

17. The apparatus of Claim 15 or 16, wherein the processor is further to determine an acoustic confidence value for the generated text result at the local device, and to apply voting to the generated and received text result if the acoustic confidence value exceeds a threshold to determine a text candidate for each word.

18. A computing system comprising:

a microphone;

a buffer of a local device to receive and record a speech utterance from the microphone;

a local speech recognition system to convert the speech utterance to a text result using a local language model including a local vocabulary in a local lexicon;

a client application to receive the text result as a command;

a network interface to send the speech utterance to at least one remote speech recognition system and to receive a text result corresponding to the utterance from the remote speech recognition system; and

a processor to compare the received text results and the local text result to determine words that are out of the local vocabulary, and to update the local language model using the out of vocabulary words.

19. The system of Claim 18, further comprising combining the generated and received text results based on the comparison to determine a final text result for the utterance and determining the out of vocabulary words using the final text result.

20. The system of Claim 18 or 19, wherein comparing comprises inserting different words from multiple received remote cloud text strings into a phone lattice and scoring the phone lattice for each combination, and substituting words from the received text result into the generated text result based on the scores.

FIG. 1





FIG. 2


FIG. 3


FIG. 4

FIG. 5





FIG. 6

A. CLASSIFICATION OF SUBJECT MATTER**G10L 15/22(2006.01)I, G10L 15/20(2006.01)I, G10L 15/02(2006.01)I, G06F 15/16(2006.01)I**

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 15/22; G10L 11/00; G10L 15/06; G10L 9/04; G10L 15/02; G06F 17/27; G10L 21/00; G10L 15/20; G06F 15/16

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: voice recognition, model, updating, lexicon, comparing, text result, and similar terms.

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2013-0132084 A1 (TIMOTHY STONEHOCKER et al.) 23 May 2013 See paragraphs [0013], [0016]–[0017], [0019]–[0021], [0028]–[0030], [0032], and [0034]–[0035]; claims 1–2; and figures 1A and 2.	1–3, 5–6, 8, 15–20
Y		4
Y	US 5,255,342 A (TSUNEO NITTA) 19 October 1993 See column 3, lines 1–22; and claim 1.	4
A	US 2014-0088967 A1 (AKINORI KAWAMURA et al.) 27 March 2014 See paragraphs [0039]–[0056] and figure 3.	1–6, 8, 15–20
A	US 2007-0276651 A1 (HARRY M. BLISS et al.) 29 November 2007 See paragraphs [0036]–[0044] and figures 3–4.	1–6, 8, 15–20
A	US 2003-0120493 A1 (SUNIL K. GUPTA) 26 June 2003 See paragraphs [0034]–[0038] and figure 2.	1–6, 8, 15–20



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

08 August 2016 (08.08.2016)

Date of mailing of the international search report

08 August 2016 (08.08.2016)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

LEE, EUN KYU

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORTInternational application No.
PCT/US2016/033623**Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)**

This international search report has not been established in respect of certain claims under Article 17(2)(a) for the following reasons:

1. ☐ Claims Nos.:
because they relate to subject matter not required to be searched by this Authority, namely:
2. ☒ Claims Nos.: 10
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
Claim 10 is unclear because it is referring to unsearchable claim 7.
3. ☒ Claims Nos.: 7,9,11-14
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

Box No. III Observations where unity of invention is lacking (Continuation of item 3 of first sheet)

This International Searching Authority found multiple inventions in this international application, as follows:

1. ☐ As all required additional search fees were timely paid by the applicant, this international search report covers all searchable claims.
2. ☐ As all searchable claims could be searched without effort justifying an additional fees, this Authority did not invite payment of any additional fees.
3. ☐ As only some of the required additional search fees were timely paid by the applicant, this international search report covers only those claims for which fees were paid, specifically claims Nos.:
4. ☐ No required additional search fees were timely paid by the applicant. Consequently, this international search report is restricted to the invention first mentioned in the claims; it is covered by claims Nos.:

Remark on Protest

- ☐ The additional search fees were accompanied by the applicant's protest and, where applicable, the payment of a protest fee.
- ☐ The additional search fees were accompanied by the applicant's protest but the applicable protest fee was not paid within the time limit specified in the invitation.
- ☐ No protest accompanied the payment of additional search fees.

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2016/033623

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2013-0132084 A1	23/05/2013	US 2015-154959 A1 US 8972263 B2 US 9330669 B2	04/06/2015 03/03/2015 03/05/2016
US 05255342 A	19/10/1993	EP 0374604 A3 EP 0374604 B1 JP 02-165388 A	02/12/1992 13/12/1995 26/06/1990
US 2014-0088967 A1	27/03/2014	JP 05706384 B2 JP 2014-063088 A US 9293137 B2	22/04/2015 10/04/2014 22/03/2016
US 2007-0276651 A1	29/11/2007	CN 101454775 A WO 2007-140047 A2 WO 2007-140047 A3	10/06/2009 06/12/2007 22/05/2008
US 2003-0120493 A1	26/06/2003	None	