



US 20150248898A1

(19) **United States**

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

(10) **Pub. No.: US 2015/0248898 A1**
(43) **Pub. Date: Sep. 3, 2015**

(54) **COMPUTER-IMPLEMENTED SYSTEMS AND METHODS FOR DETERMINING AN INTELLIGIBILITY SCORE FOR SPEECH**

(52) **U.S. CL.**
CPC *G10L 25/60* (2013.01); *G10L 15/187* (2013.01); *G10L 15/02* (2013.01)

(71) Applicant: **Educational Testing Service**, Princeton, NJ (US)

(57) **ABSTRACT**

(72) Inventors: **Anastassia Loukina**, Bellport, NY (US);
Keelan Evanini, Pennington, NJ (US)

(21) Appl. No.: **14/632,231**

(22) Filed: **Feb. 26, 2015**


Related U.S. Application Data

(60) Provisional application No. 61/945,856, filed on Feb. 28, 2014.

Publication Classification

(51) **Int. Cl.**
G10L 25/60 (2006.01)
G10L 15/02 (2006.01)
G10L 15/187 (2006.01)

Systems and methods are provided for generating an intelligibility score for speech of a non-native speaker. Words in a speech recording are identified using an automated speech recognizer, where the automated speech recognizer provides a string of words identified in the speech recording, and where the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words. For a particular word in the string of words, a context metric value is determined based upon a usage of the particular word within the string of words. An acoustic score for the particular word is determined based on the acoustic model likelihood score for the particular word from the automated speech recognizer. An intelligibility score is determined for the particular word based on the acoustic score for the particular word and the context metric value for the particular word.






Fig. 1




Fig. 2

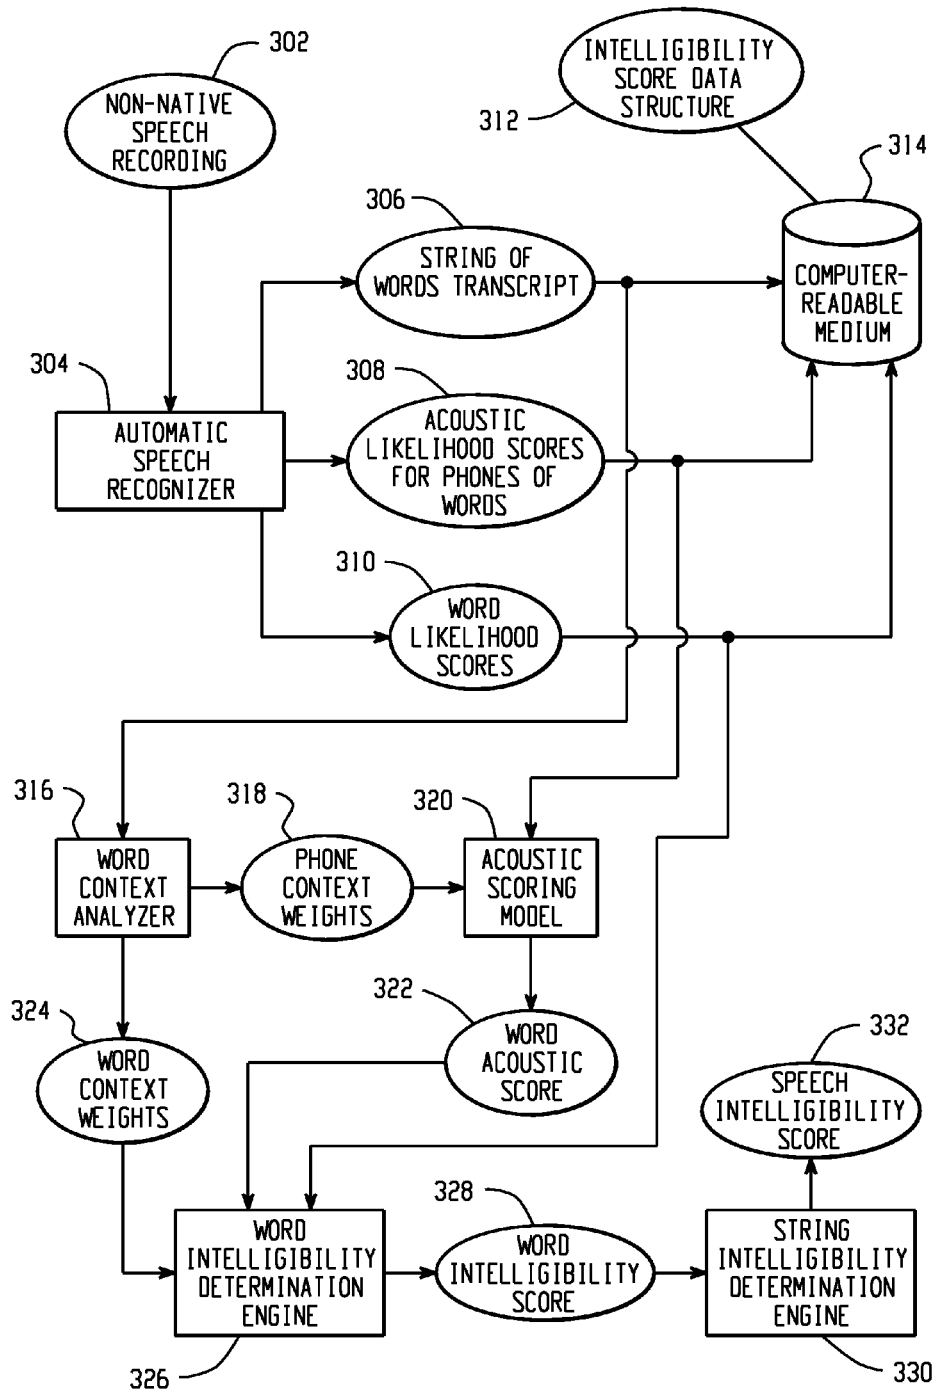
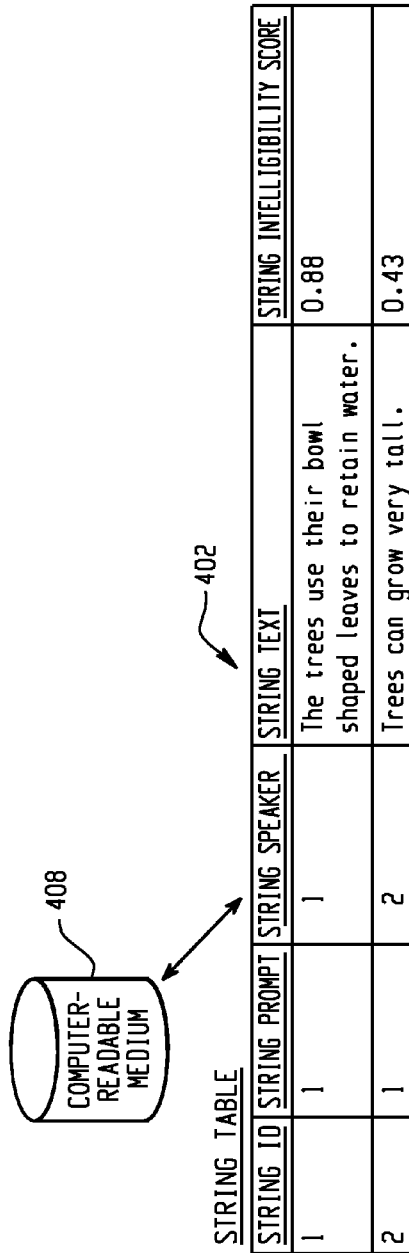


Fig. 3



408

COMPUTER-READABLE MEDIUM

402

STRING ID	STRING PROMPT	STRING SPEAKER	STRING TEXT	STRING INTELLIGIBILITY SCORE
1		1	The trees use their bowl shaped leaves to retain water.	0.88
2		2	Trees can grow very tall.	0.43

404

WORD ID	STRING ID	WORD	WORD POSITION	WORD CONTEXT	WORD ACOUSTIC SCORE	WORD INTELLIGIBILITY
273	1	their	4	0.4	0.34	0.74
274	1	bowl	5	0.4	0.62	1.02
275	1	shaped	6	0.3	0.9	1.2

406

PHONE ID	WORD ID	PHONE	PHONE POSITION	PHONE CONTEXT WEIGHT	PHONE ACOUSTIC SUB-SCORE
1024	274	b	1	1.5	0.7
1025	274	ow	2	0.6	0.9
1026	274	l	3	0.9	0.3

Fig. 4


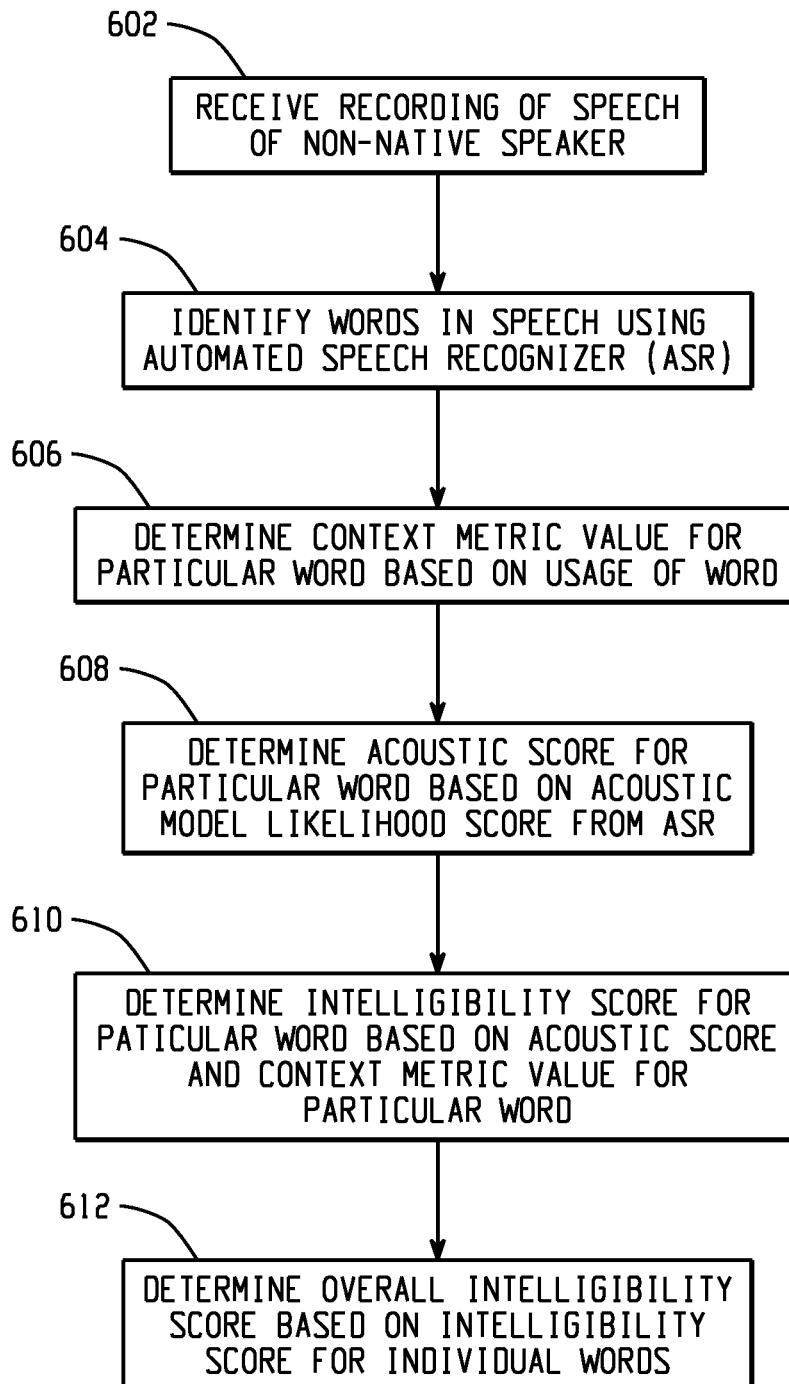


Fig. 5

*Fig. 6*

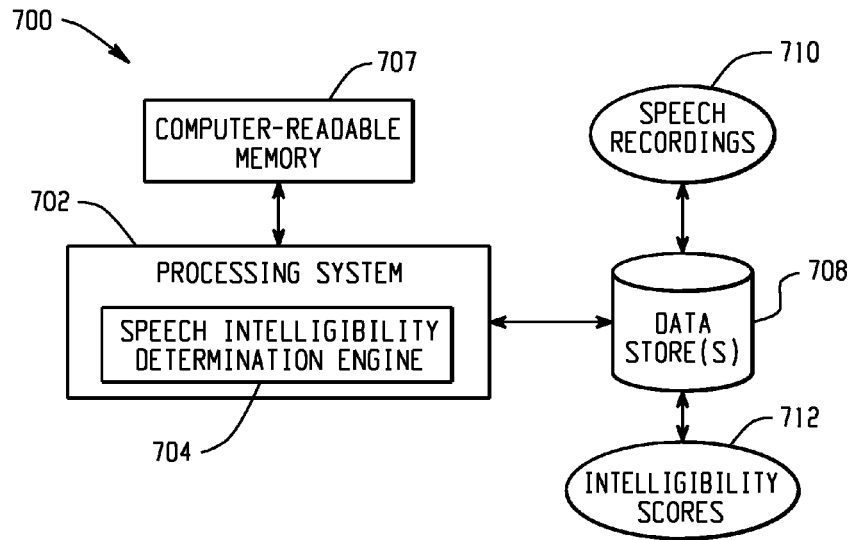


Fig. 7A

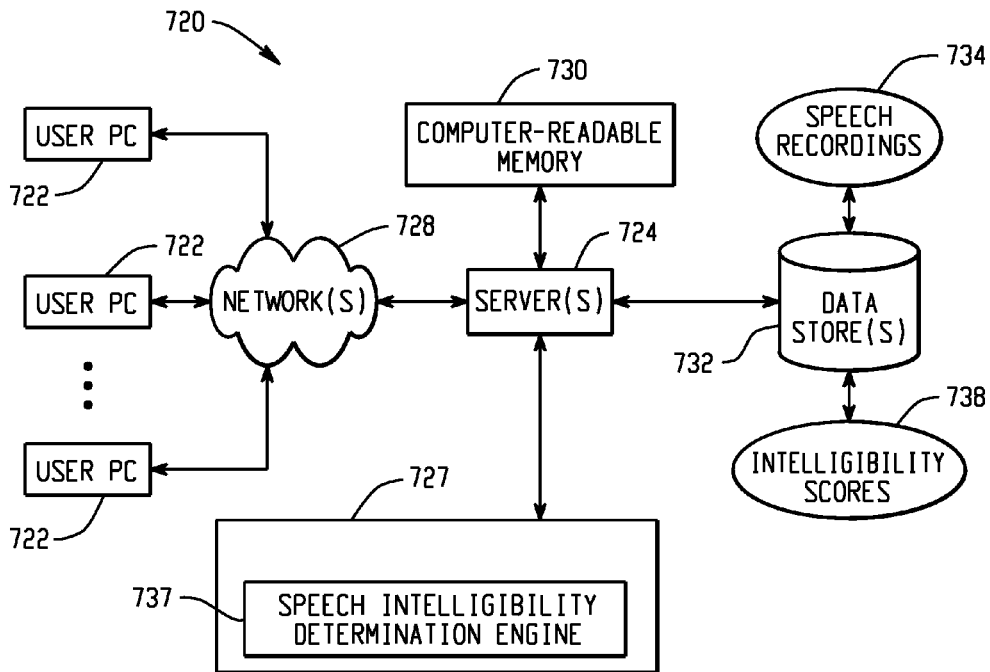


Fig. 7B




Fig. 7C

COMPUTER-IMPLEMENTED SYSTEMS AND METHODS FOR DETERMINING AN INTELLIGIBILITY SCORE FOR SPEECH

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to U.S. Provisional Patent Application No. 61/945,856, filed Feb. 28, 2014, entitled “A Method to Compute Intelligibility Scores, Identify Mispronounced Words and Evaluate Accuracy of Spoken Responses for Automated Speech Scoring,” which is incorporated herein by reference in its entirety.

FIELD

[0002] The technology described in this patent document relates generally to speech processing and scoring and more particularly to automatically scoring the intelligibility of recorded speech.

BACKGROUND

[0003] The ability to communicate well is a crucial skill that greatly improves quality of life for a person. Pronunciation evaluation and feedback can be important tools for improving a speaker’s ability to be easily understood. For example, a speaker learning a second language can improve their ability to be understood by native speakers of that second language by improving their pronunciation, such that it is more similar to a native speaker of that second language. Traditionally, speech pronunciation has been evaluated by comparing an examinee’s speech to that of a native speaker. For example, a human evaluator can listen to the speaker and identify differences in the speaker’s pronunciation compared to a native speaker. The evaluator then provides a score and/or feedback to the examinee, which can aid the examinee in improving their communication ability.

SUMMARY

[0004] Systems and methods are provided for generating an intelligibility score for speech of a non-native speaker. Words in a speech recording are identified using an automated speech recognizer, where the automated speech recognizer provides a string of words identified in the speech recording, and where the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words. For a particular word in the string of words, a context metric value is determined based upon a usage of the particular word within the string of words. An acoustic score for the particular word is determined based on the acoustic model likelihood score for the particular word from the automated speech recognizer. An intelligibility score is determined for the particular word based on the acoustic score for the particular word and the context metric value for the particular word, and an overall intelligibility score for the string of words is determined based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.

[0005] As another example, a computer-implemented system for generating an intelligibility score for speech of a non-native speaker includes a processing system and a non-transitory computer-readable medium encoded to contain instructions for commanding the execute steps of a method. In the method, words in a speech recording are identified using an automated speech recognizer, where the automated

speech recognizer provides a string of words identified in the speech recording, and where the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words. For a particular word in the string of words, a context metric value is determined based upon a usage of the particular word within the string of words. An acoustic score for the particular word is determined based on the acoustic model likelihood score for the particular word from the automated speech recognizer. An intelligibility score is determined for the particular word based on the acoustic score for the particular word and the context metric value for the particular word, and an overall intelligibility score for the string of words is determined based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.

[0006] As a further example, a non-transitory computer-readable medium is encoded with instructions for commanding a processing system to execute a method of generating an intelligibility score for speech of a non-native speaker. In the method, words in a speech recording are identified using an automated speech recognizer, where the automated speech recognizer provides a string of words identified in the speech recording, and where the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words. For a particular word in the string of words, a context metric value is determined based upon a usage of the particular word within the string of words. An acoustic score for the particular word is determined based on the acoustic model likelihood score for the particular word from the automated speech recognizer. An intelligibility score is determined for the particular word based on the acoustic score for the particular word and the context metric value for the particular word, and an overall intelligibility score for the string of words is determined based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] FIG. 1 is a block diagram depicting a computer implemented speech intelligibility determination engine.

[0008] FIG. 2 is a diagram depicting example components of a speech intelligibility determination engine.

[0009] FIG. 3 is a data flow diagram depicting example operations of a speech intelligibility determination engine.

[0010] FIG. 4 is a diagram depicting an example intelligibility score data structure stored in a computer-readable storage medium.

[0011] FIG. 5 is a diagram depicting four example intelligibility feedback displays.

[0012] FIG. 6 is a flow diagram depicting a computer-implemented method for generating an intelligibility score for speech of a non-native speaker.

[0013] FIGS. 7A, 7B, and 7C depict example systems for implementing the approaches described herein for generating an intelligibility score for speech.

DETAILED DESCRIPTION

[0014] FIG. 1 is a block diagram depicting a computer implemented speech intelligibility determination engine. Traditionally, pronunciation evaluation has been used as a proxy for measuring how well an examinee can be understood. Differences between the examinee’s pronunciation and a typical native speaker’s pronunciation are identified,

and a pronunciation score is provided that is an approximation of the intelligibility of that examinee's speech. While a straight, phone-by-phone or word-by-word pronunciation comparison does provide some insight into speech intelligibility, it is not a direct measure of intelligibility. This is because a pronunciation-only evaluation ignores the context of the phones or words present in the speech. That is, while almost every pronunciation error affects the intelligibility of speech, certain errors affect intelligibility more than others based when those errors occur in the speech (e.g., mispronunciation in combination with a poor word choice is more detrimental to intelligibility than mispronunciation of a correct word in a context).

[0015] Context can provide significant clues to a listener as to sounds or words presented in speech, such that a listener may be able to understand speech or portions of speech even when the speaker's pronunciation is sub-optimal. For example, when a listener is conversant in a topic discussed by a speaker, the listener can more easily understand what is being spoken based on an understanding of the context of the speech despite pronunciation errors. In contrast, it can be much more difficult to understand a speaker making pronunciation errors when the discussion is based on a topic of little knowledge to the listener. Thus, while even a very proficient non-native speaker's pronunciation may differ from a native speaker's (e.g., an accent), the context of occurrences of those differences can highly affect the intelligibility or understandability of that speech.

[0016] Systems and methods herein provide mechanisms for measuring the intelligibility of speech by considering context of speech (e.g., phones, words within the speech) in addition to pronunciation. With reference to FIG. 1, a computer implemented speech intelligibility determination engine 102 receives a speech sample to analyze, such as a recording of speech 104 by a non-native speaker. The speech intelligibility determination engine 102 analyzes the recording 104 based on pronunciation of phones and words within the speech in combination with the context of those phones and words to generate a speech intelligibility score 106.

[0017] FIG. 2 is a diagram depicting example components of a speech intelligibility determination engine. A speech intelligibility determination engine 202 receives speech to be analyzed, such as a non-native speech recording 204. An automatic speech recognizer 206 is configured to receive the non-native speech recording 204 and to recognize words within the recording 204 and certain metrics of the speech in the recording 204 (e.g., via a forced alignment procedure). Those words and metrics can be stored for downstream usage, such as in an intelligibility score data structure 208 on a non-transitory computer-readable medium 210. In one example, the automatic speech recognizer 206 is configured to identify a string of words identified in the speech recording 204. The automatic speech recognizer is further configured to provide an acoustic model likelihood score for each word identified in the speech recording 204. In one embodiment, the automatic speech recognizer identifies multiple metrics for each word in the recording. In such an example, the automatic speech recognizer provides an acoustic score for each phone in the word that indicates a quality of pronunciation of that phone (e.g., for the word "bowl," an acoustic score is provided for the "b," the "ow," and the "l" sounds). The automatic speech recognizer 206 further provides a confidence score for the word (e.g., a score indicating how confident the recognizer 206 is that "bowl" is the word spoken in

the recording 204). In one example, the automatic speech recognizer or other entities provide other metrics, such as a speech rate metric indicating a speed at which the speaker in the recording 204 is speaking or a predictability of a word in a sentence based on other words in the sentence.

[0018] The speech intelligibility determination engine 202 further includes a word context analyzer 212 configured to provide metrics, such as for storage in the intelligibility score data structure 208, that indicate context of phones and/or words detected in the speech 204 by the automatic speech recognizer 206. In one embodiment, the word context analyzer 212 identifies phone-level factors for phones in a word, such as position of the phone in the word and whether the phone is stressed in the word. The word context analyzer 212 is further configured to provide word-level factors such as part of speech and lexical frequency. The phone and/or word context metrics are utilized by downstream entities, such as to provide weights to detected pronunciation errors to measure an effect of those pronunciation errors on the intelligibility of the speech. For example, pronunciation errors at the end of a word or in a non-accented syllable may be weighted less than pronunciation errors at the beginning of a word or in an accented syllable. Further, pronunciation errors of uncommon words in a sentence may be weighted more heavily than pronunciation of common words or words that are likely to occur in a sentence based on other words in the sentence (e.g., mispronunciation of the word "chicken" is weighted more heavily in the phrase "The dean handed him a chicken at graduation;" compared to weighting for the phrase "We ate fried chicken for dinner.").

[0019] The speech intelligibility determination engine 202 further includes components (e.g., software modules) for scoring the speech recording 204. In one embodiment, an acoustic scoring model 214 is configured to analyze the pronunciation of phones within words along with the context of those phones to determine word acoustic scores for those words. Further, a word intelligibility determination engine 216 is configured to analyze word-level pronunciation and word context to score intelligibility at the word level. Phone-level scores from the acoustic scoring model 214 and word-level scores from the word intelligibility determination engine 216 are utilized for some or all of the words in the recording to generate a speech intelligibility score 218, such as at the word intelligibility determination engine 216.

[0020] FIG. 3 is a data flow diagram depicting example operations of a speech intelligibility determination engine. A non-native speech recording 302 is received by an automatic speech recognizer 304. In one example, the automatic speech recognizer 304 is trained using samples of native and non-native speech using acoustic model training procedures. In other examples, the recognizer 304 is trained using native-only or non-native-only speakers. The automatic speech recognizer provides a transcript 306 of a string of words recognized in the recording, acoustic likelihood scores 308 indicating confidence or quality of pronunciation of individual phones or words in the recording 302, and word likelihood scores indicating a confidence of the automatic speech recognizer 304 that it has correctly recognized those words in the recording 302. The transcript 306 and the metrics 308, 310 are stored in an intelligibility score data structure 312 in a computer readable medium 314. In an example, the acoustic likelihood scores 308 are computed as average raw values of acoustic model likelihood scores for each phone or more complex measures such as goodness of pronunciation scores

based on acoustic model likelihood scores and prior probabilities for each phone. Word likelihood scores 310 are provided based on a confidence of the automatic speech recognizer 304 that it has identified the correct word in the transcript 306.

[0021] A word context analyzer 316 receives the transcript 306 and utilizes that transcript to determine one or more context metrics for words within the transcript 306. In one embodiment, the word context analyzer 316 determines phone-level context metrics for phones in words of the transcript 306, where in the example of FIG. 3, those phone-level context metrics (e.g., metrics based on the context of a particular phone within a particular word) are utilized to generate phone context weights 318 that are transmitted to an acoustic scoring model 320. Phone context weights 318 can be based on a variety of phone-level context metrics including lexical stress, type of phone (e.g., vowel or consonant), position in word, position in syllable, and syllable weight. In one example, the phone context weights can be calculated by a computer processing system by referencing a computer-readable storage medium that includes a table of weights based on phone-level context metric values (e.g., increase context weight 0.1 for vowels, increase context weight 0.2 for phones at the start of a word).

[0022] The acoustic scoring model 320 receives acoustic likelihood scores 308 for each phone in the transcript 306 as well as the phone context weights 318 for those phones from the word context analyzer 316. The acoustic scoring model 320 transforms those values 308, 318 into a word acoustic score 322 for each word in the transcript 306. In one example, the acoustic scoring model is a computerized model that is configured to generate a word acoustic score 322 for a particular word based on a per-phone average of the product of a phone context weight 318 and acoustic likelihood score 308 for each phone of the particular word.

[0023] The word context analyzer 316 is further tasked with determining word-level context metrics for words in the transcript 306. For example, a word-level context metric may identify a likelihood of a particular word in the transcript 306 appearing at its position in the transcript 306 based on other words in the transcript 306 as analyzed with reference to a corpus of training documents. As another example, a word-level context metric may identify how commonly the particular word is used in a corpus of documents. Other word-level context metrics can be based on part of speech, position in sentence, grammar errors, and position in phrase. The word-level context metrics are utilized to determine word context weights 324 that are provided to a word intelligibility determination engine 326.

[0024] The word intelligibility determination engine 326 receives word likelihood scores 310 for each word in the transcript 306 from the automatic speech recognizer 304 as well as the word context weights 324 for those words from the word context analyzer. The word intelligibility determination engine 326 transforms those values 310, 324 into a word level score for each word in the transcript 306. In one example, a word level score for a particular word is based on a product of the word likelihood score 310 and the word context weight 324 the particular word.

[0025] In the example of FIG. 3, the word intelligibility determination engine 326 transforms the word acoustic score 322 for the particular word and the word level score for the particular word into a word intelligibility score 328. The word intelligibility score 328 is received by a string intelligibility

determination engine 330 that generates a speech intelligibility score 332 for the speech recording. In one embodiment, the speech intelligibility score 332 is based on an average of word intelligibility scores 328 for the recording or other combination of those word intelligibility scores 328.

[0026] Following is an example determination of a word intelligibility score determination process. In this example, the automatic speech recognizer 304 recognizes the sentence: "The trees use their bowl-shaped leaves to retain water." This transcript 306 is stored in the intelligibility score data structure 312 along with various metrics and intermediate scores used in determining a speech intelligibility score 322 for the phrase. For the particular word "bowl," the word intelligibility score 328 is calculated according to:

$$IS_{bowl} = f(AM_{bowl}, LM_{bowl}, CS_{bowl}, Lex_{bowl}), \quad (\text{eq. 1})$$

where IS_{bowl} is the intelligibility score 328 for the word "bowl," AM_{bowl} is a word acoustic score 322 for the word "bowl," LM_{bowl} is the probability of the word "bowl" in a given context using different lengths of left and right context (e.g., the probability of "bowl" in the following phrases: "their bowl," "use their bowl," "bowl-shaped," "bowl-shaped leaves," etc.), CS_{bowl} is the automatic speech recognizer 304 confidence score outputted as a word likelihood score 310, and Lex_{bowl} is a lexical score for the word bowl based on a combination of measurements such as lexical frequency in one or more corpora and point-wise mutual information between "bowl" and other words in the sentence.

[0027] In one example, the word acoustic score (AM_{bowl}) is calculated according to

$$AM_{bowl} = AM_b * w_{initial} * w_{consonant} + AM_{ow} * w_{vowel} * w_{stressed} + AM_1 + \text{speech_rate}_{bowl} \quad (\text{eq. 2})$$

where AM_b is an acoustic likelihood score 308 for the "b" phone in "bowl," $w_{initial}$ is a weight associated with "b" being at the beginning of the word, $w_{consonant}$ is a weight associated with the "b" phone being a consonant, AM_{ow} is an acoustic likelihood score 308 for the "ow" phone, w_{vowel} is a weight associated with "ow" being a vowel phone, $w_{stressed}$ is a weight associated with "ow" being an accented phone, AM_1 is an acoustic likelihood score 308 for the "1" phone, and $\text{speech_rate}_{bowl}$ is a metric associated with the rate of speech of the word "bowl."

[0028] Having calculated the component terms of eq. 1, the word intelligibility determination engine 326 determines a word intelligibility score 328 for the word "bowl," according to a formula such as:

$$IS_{bowl} = a + b_1 * AM_{bowl} + b_2 * LM_{bowl} + b_3 * CS_{bowl} + b_4 * Lex_{bowl} \quad (\text{eq. 3})$$

where a , b_1 , b_2 , b_3 , and b_4 are constant weights determined via a model training operation, such as a model training operation that analyzes correlations between metrics of recognized training speech and human-provided intelligibility scores for that training speech.

[0029] FIG. 4 is a diagram depicting an example intelligibility score data structure stored in a computer-readable storage medium. The example data structure of FIG. 4 is stored in three tables 402, 404, 406 of a relational database on the computer-readable medium 408. A string table 402 stores recognized text of strings as well as data indicating a prompt with which the string is associated, a speaker of the string, and a string intelligibility score for the string. A word table 404 is configured to contain data associated with words in a string. In the example of FIG. 4, the word table 404 contains data

associated with three words of string ID 1. The word table identifies a position of a word in its string, a word context weight for that word, a word acoustic score for that word, and a word intelligibility score for that word. A phone table 406 includes data associated with phones of individual words. The phone table includes data for phones of the word "bowl" including a position of phones within that word, phone context weights for those phones, and phone associated sub-scores (e.g., a product of a phone context weight and a acoustic likelihood score for that phone).

[0030] In addition to outputting an overall speech intelligibility score 322 for a speech sample, a speech intelligibility determination engine can be configured to provide feedback to a speaker, such as via a graphical user interface. FIG. 5 is a diagram depicting four example intelligibility feedback displays. At 502, a first display is provided that highlights words having word intelligibility scores less than a threshold value. At 504, a second display uses highlighting and question marks to identify words having word intelligibility scores less than a threshold value. In a third example at 506, different color highlights or a gradient of colors are used to indicate word intelligibility scores. In one embodiment, well spoken words having intelligibility scores above a first threshold are highlighted in green, poorly spoken words having intelligibility scores below a second threshold are highlighted in red, and other words having intelligibility scores between the first and second thresholds are highlighted in yellow. At 508, an example is provided that depicts colored highlighting on a timeline, indicating portions of a speech timeline where intelligibility scores fell within different ranges demarcated by the first and second thresholds.

[0031] FIG. 6 is a flow diagram depicting a computer-implemented method for generating an intelligibility score for speech of a non-native speaker. A speech recording of a non-native speaker is received at 602. At 604, words in a speech recording are identified using an automated speech recognizer, where the automated speech recognizer provides a string of words identified in the speech recording, and where the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words. For a particular word in the string of words, at 606, a context metric value is determined based upon a usage of the particular word within the string of words. An acoustic score for the particular word is determined at 608 based on the acoustic model likelihood score for the particular word from the automated speech recognizer. At 610, an intelligibility score is determined for the particular word based on the acoustic score for the particular word and the context metric value for the particular word, and an overall intelligibility score for the string of words is determined at 612 based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.

[0032] FIGS. 7A, 7B, and 7C depict example systems for implementing the approaches described herein for generating an intelligibility score for speech. For example, FIG. 7A depicts an exemplary system 700 that includes a standalone computer architecture where a processing system 702 (e.g., one or more computer processors located in a given computer or in multiple computers that may be separate and distinct from one another) includes a speech intelligibility determination engine 704 being executed on the processing system 702. The processing system 702 has access to a computer-readable memory 707 in addition to one or more data stores 708. The one or more data stores 708 may include speech

recordings 710 as well as intelligibility scores 712. The processing system 702 may be a distributed parallel computing environment, which may be used to handle very large-scale data sets.

[0033] FIG. 7B depicts a system 720 that includes a client-server architecture. One or more user PCs 722 access one or more servers 724 running a speech intelligibility determination engine 737 on a processing system 727 via one or more networks 728. The one or more servers 724 may access a computer-readable memory 730 as well as one or more data stores 732. The one or more data stores 732 may include speech recordings 734 as well as intelligibility scores 738.

[0034] FIG. 7C shows a block diagram of exemplary hardware for a standalone computer architecture 750, such as the architecture depicted in FIG. 7A that may be used to include and/or implement the program instructions of system embodiments of the present disclosure. A bus 752 may serve as the information highway interconnecting the other illustrated components of the hardware. A processing system 754 labeled CPU (central processing unit) (e.g., one or more computer processors at a given computer or at multiple computers), may perform calculations and logic operations required to execute a program. A non-transitory processor-readable storage medium, such as read only memory (ROM) 758 and random access memory (RAM) 759, may be in communication with the processing system 754 and may include one or more programming instructions for performing the method of generating an intelligibility score for speech. Optionally, program instructions may be stored on a non-transitory computer-readable storage medium such as a magnetic disk, optical disk, recordable memory device, flash memory, or other physical storage medium.

[0035] In FIGS. 7A, 7B, and 7C, computer readable memories 707, 730, 758, 759 or data stores 708, 732, 783, 784, 788 may include one or more data structures for storing and associating various data used in the example systems for generating an intelligibility score for speech. For example, a data structure stored in any of the aforementioned locations may be used to store data from XML files, initial parameters, and/or data for other variables described herein. A disk controller 790 interfaces one or more optional disk drives to the system bus 752. These disk drives may be external or internal floppy disk drives such as 783, external or internal CD-ROM, CD-R, CD-RW or DVD drives such as 784, or external or internal hard drives 785. As indicated previously, these various disk drives and disk controllers are optional devices.

[0036] Each of the element managers, real-time data buffer, conveyors, file input processor, database index shared access memory loader, reference data buffer and data managers may include a software application stored in one or more of the disk drives connected to the disk controller 790, the ROM 758 and/or the RAM 759. The processor 754 may access one or more components as required.

[0037] A display interface 787 may permit information from the bus 752 to be displayed on a display 780 in audio, graphic, or alphanumeric format. Communication with external devices may optionally occur using various communication ports 782.

[0038] In addition to these computer-type components, the hardware may also include data input devices, such as a keyboard 779, or other input device 781, such as a microphone, remote control, pointer, mouse and/or joystick.

[0039] Additionally, the methods and systems described herein may be implemented on many different types of pro-

cessing devices by program code comprising program instructions that are executable by the device processing subsystem. The software program instructions may include source code, object code, machine code, or any other stored data that is operable to cause a processing system to perform the methods and operations described herein and may be provided in any suitable language such as C, C++, JAVA, for example, or any other suitable programming language. Other implementations may also be used, however, such as firmware or even appropriately designed hardware configured to carry out the methods and systems described herein.

[0040] The systems' and methods' data (e.g., associations, mappings, data input, data output, intermediate data results, final data results, etc.) may be stored and implemented in one or more different types of computer-implemented data stores, such as different types of storage devices and programming constructs (e.g., RAM, ROM, Flash memory, flat files, databases, programming data structures, programming variables, IF-THEN (or similar type) statement constructs, etc.). It is noted that data structures describe formats for use in organizing and storing data in databases, programs, memory, or other computer-readable media for use by a computer program.

[0041] The computer components, software modules, functions, data stores and data structures described herein may be connected directly or indirectly to each other in order to allow the flow of data needed for their operations. It is also noted that a module or processor includes but is not limited to a unit of code that performs a software operation, and can be implemented for example as a subroutine unit of code, or as a software function unit of code, or as an object (as in an object-oriented paradigm), or as an applet, or in a computer script language, or as another type of computer code. The software components and/or functionality may be located on a single computer or distributed across multiple computers depending upon the situation at hand.

[0042] The computerized approaches described herein, which utilize, e.g., various computer models trained according to sample data, are very different from conventional human scoring of the quality of speech of a non-native language speaker. In conventional human scoring of the quality of speech of a non-native language speaker, a human grader listens to the speech of the non-native speaker and makes a holistic judgment about the quality of the speech and assigns a score. Conventional human grading of the quality of speech of a non-native speaker does not involve the use of the computer models, associated variables, training of the models based on sample data to calculate weights of various features or variables, computer processing to parse speech samples to be scored and representing such parsed speech samples with suitable data structures, and applying the computer models to those data structures to score the speech samples, as described herein.

[0043] While the disclosure has been described in detail and with reference to specific embodiments thereof, it will be apparent to one skilled in the art that various changes and modifications can be made therein without departing from the spirit and scope of the embodiments. Thus, it is intended that the present disclosure cover the modifications and variations of this disclosure provided they come within the scope of the appended claims and their equivalents.

It is claimed:

1. A computer-implemented method of generating an intelligibility score for speech of a non-native speaker, comprising:

receiving a recording of speech of a non-native speaker at a processing system;

identifying words in the speech recording using a computerized automated speech recognizer, wherein the automated speech recognizer provides a string of words identified in the speech recording based on a computerized acoustic model, and wherein the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words;

for a particular word in the string of words, determining a context metric value with the processing system based upon a usage of the particular word within the string of words;

determining an acoustic score with the processing system for the particular word based on the acoustic model likelihood score for the particular word from the automated speech recognizer;

determining an intelligibility score with the processing system for the particular word based on the acoustic score for the particular word and the context metric value for the particular word; and

determining an overall intelligibility score with the processing system for the string of words based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.

2. The method of claim 1, wherein the acoustic model likelihood score for the particular word is based on acoustic model likelihood sub-scores for each phone of the particular word.

3. The method of claim 2, further comprising determining a phone context weighting value for each phone of the particular word.

4. The method of claim 3, wherein the phone context weighting value for a particular phone is based on one or more of a position of the particular phone in the word, whether the particular phone is a vowel phone, and whether the particular phone is stressed.

5. The method of claim 3, wherein the acoustic model likelihood score for the particular word is based on a sum of the acoustic model likelihood sub-score multiplied by the phone context weighting value for each phone of the particular word.

6. The method of claim 3, wherein the acoustic model likelihood score for the particular word is calculated according to:

$$AM_{word} = \sum_{i=1}^N AM_i * w_i,$$

where AM_{word} is the acoustic model likelihood score for the particular word, N is the number of phones in the particular word, AM_i is the acoustic model likelihood sub-score for the i th phone of the particular word, and w_i is the phone context weighting value for the i th phone of the particular word.

7. The method of claim 1, wherein the context metric value is based on a probability of the particular word being used with the other words in the string of words, a part of speech of the particular word, and a frequency of occurrence of the particular word in a corpus of documents.

8. The method of claim 1, further comprising:
 comparing the intelligibility score for the particular word to an intelligibility threshold to determine whether the particular word meets an intelligibility criterion;
 wherein the overall intelligibility score for the string of words is based on a proportion of words in the string of words that meet the intelligibility criterion.
9. The method of claim 8, wherein the intelligibility threshold is set based on a model training operation that utilizes training recordings of speech manually scored for intelligibility by one or more human scorers.
10. The method of claim 9, wherein the model training operation comprises determining training intelligibility scores for words within the training recordings of speech and correlating the training intelligibility scores with the manual scores assigned by the one or more human scorers.
11. The method of claim 10, wherein correlating comprises establishing a function which maps intelligibility scores to manual scores.
12. The method of claim 1, wherein the overall intelligibility score is based on an average intelligibility score of all words in the string of words or a subset of words considered eligible for computing the intelligibility score.
13. The method of claim 1, wherein the words in the speech recording are identified by the automated speech recognizer.
14. The method of claim 1, wherein a second word having a same acoustic score as the acoustic score for the particular word is determined to have a different intelligibility score based on differing context metric values between the particular word and the second word.
15. The method of claim 1, wherein the automated speech recognizer further provides an acoustic model likelihood score for each phone within each word, a language model likelihood score for each word, and a confidence score for each word.
16. The method of claim 1, wherein the context metric value is further based upon a part of speech of the particular word or a lexical frequency of the particular word.
17. The method of claim 1, further comprising providing a display that includes the identified words and the overall intelligibility score.
18. A computer-implemented system for generating an intelligibility score for speech of a non-native speaker, comprising:
 a processing system;
 a non-transitory computer-readable medium encoded to contain:
 a recording of speech of a non-native speaker;
 an intelligibility score data structure comprising records associated with each word of a string of words, wherein a record for a particular word in the string of words includes fields for storing an acoustic score for the particular word, a context metric value for the particular word, and an intelligibility score for the particular word;
 instructions for commanding the processing system to execute steps comprising:
 identifying words in the speech recording using an automated speech recognizer, wherein the automated speech recognizer provides a string of words identified in the speech recording, and wherein the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words;
 for the particular word, determining a context metric value based upon a usage of the particular word within the string of words and storing the context metric value in the intelligibility score data structure record for the particular word;
 determining an acoustic score for the particular word based on the acoustic model likelihood score for the particular word from the automated speech recognizer and storing the acoustic score in the intelligibility score data structure record for the particular word;
 determining an intelligibility score for the particular word based on the acoustic score for the particular word and the context metric value for the particular word and storing the intelligibility score in the intelligibility score data structure record for the particular word; and
 determining an overall intelligibility score for the string of words based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.
19. The system of claim 18, wherein the acoustic model likelihood score for the particular word is based on acoustic model likelihood sub-scores for each phone of the particular word.
20. The system of claim 19, wherein the steps further comprise determining a phone context weighting value for each phone of the particular word.
21. The system of claim 20, wherein the phone context weighting value for a particular phone is based on one or more of a position of the particular phone in the word, whether the particular phone is a vowel phone, and whether the particular phone is stressed.
22. The system of claim 20, wherein the acoustic model likelihood score for the particular word is based on a sum of the acoustic model likelihood sub-score multiplied by the phone context weighting value for each phone of the particular word.
23. A non-transitory computer-readable medium encoded with instructions for commanding a processing system to execute a method of generating an intelligibility score for speech of a non-native speaker, comprising:
 receiving a recording of speech of a non-native speaker;
 identifying words in the speech recording using an automated speech recognizer, wherein the automated speech recognizer provides a string of words identified in the speech recording, and wherein the automated speech recognizer further provides an acoustic model likelihood score for each word in the string of words;
 for a particular word in the string of words, determining a context metric value based upon a usage of the particular word within the string of words;
 determining an acoustic score for the particular word based on the acoustic model likelihood score for the particular word from the automated speech recognizer;
 determining an intelligibility score for the particular word based on the acoustic score for the particular word and the context metric value for the particular word; and
 determining an overall intelligibility score for the string of words based on the intelligibility score for the particular word and intelligibility scores for other words in the string of words.