



(51) International Patent Classification:
G10L 15/00 (2013.01)

(21) International Application Number:
PCT/US2018/021670

(22) International Filing Date:
09 March 2018 (09.03.2018)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/470,910 14 March 2017 (14.03.2017) US
15/706,128 15 September 2017 (15.09.2017) US

(63) Related by continuation (CON) or continuation-in-part (CIP) to earlier application:
US 15/706,128 (CON)
Filed on 15 September 2017 (15.09.2017)

(71) Applicant: TEXAS INSTRUMENTS INCORPORATED [US/US]; P.O. Box 655474, Mail Station 3999, Dallas, TX 75265-5474 (US).

(71) Applicant (for JP only): TEXAS INSTRUMENTS JAPAN LIMITED [JP/JP]; 24-1, Nishi-shinjuku 6-chome, Shinjuku-ku Tokyo, 160-8366 (JP).

(72) Inventors: PANDE, Tarkesh; 3807 Marchwood Drive, Richardson, TX 75082 (US). NETSCH, Lorin, Paul; 1100 Candlewood Drive, Allen, TX 75002 (US). MAGEE, David, Patrick; 1108 Drum Court, Allen, TX 75013 (US).

(74) Agent: DAVIS, Michael, A., Jr. et al.; Texas Instruments Incorporated, P.O. Box 655474, Mail Station 3999, Dallas, TX 75265-5474 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, 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,

(54) Title: QUALITY FEEDBACK ON USER-RECORDED KEYWORDS FOR AUTOMATIC SPEECH RECOGNITION SYSTEMS




FIG. 2

(57) Abstract: In an automated speech recognition system (50), a microphone (52) records a keyword spoken by a user. A front end (62) divides the recorded keyword into a plurality of subunits, each containing a segment of recorded audio, and extracts a set of features from each of the plurality of subunits. A decoder (64) assigns one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit. A quality evaluation component (66) calculates a score representing a quality of the keyword from the content classes assigned to the plurality of subunits.

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

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

QUALITY FEEDBACK ON USER-RECORDED KEYWORDS FOR AUTOMATIC SPEECH RECOGNITION SYSTEMS

[0001] This relates generally to automatic speech recognition systems, and more particularly to providing quality feedback on user-recorded keywords for automatic speech recognition systems.

BACKGROUND

[0002] The internet of things (IOT) is at its infancy and is beginning to find more applications in homes, building automation and industrial settings. It is primarily driven by the ability to have chipsets that are able to run complicated algorithms at low-power and at low cost. A key enabling technology is the human-machine interaction via voice commands. One example of this scenario is the ability of devices to recognize and respond to short voice commands. These are known as keyword spotting applications.

SUMMARY

[0003] In one example of an automated speech recognition system, a microphone records a keyword spoken by a user. A front end that divides the recorded keyword into a plurality of subunits, each containing a segment of recorded audio, and extracts a set of features from each of the plurality of subunits. A decoder assigns one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit. A quality evaluation component calculates a score representing a quality of the keyword from the content classes assigned to the plurality of subunits.

[0004] In another example of a method, an input from a user is received at an automatic speech recognition system indicating that the user intends to submit a new keyword to the system, the keyword spoken by the user is recorded. A score is assigned to the keyword indicative of a quality of the keyword. The quality of the keyword represents at least one of a recognition rate of the keyword at a decoder of the automatic speech recognition system and a false positive rate of the keyword at the decoder of the automatic speech recognition system. Substantially real-time feedback representing the assigned score is provided to the user.

[0005] In a further example of an automated speech recognition system, a microphone records a keyword spoken by a user, and a front end divides the recorded keyword into a plurality of

subunits, each containing a segment of recorded audio, and extracts a set of features from each of the plurality of subunits. A decoder assigns one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit. A quality evaluation component calculates a score representing a quality of the keyword from the content classes assigned to the plurality of subunits. An output device that provides feedback on the quality of the keyword to the user. Each of the front end, the decoder, and the quality evaluation component are configured such that the output device provides the feedback substantially in real time.

BRIEF DESCRIPTION OF THE DRAWINGS

[0006] FIG. 1 illustrates one example of a system for evaluating the quality of a keyword for an automated speech recognition system;

[0007] FIG. 2 illustrates another example of a system for evaluating the quality of a keyword for an automated speech recognition system;

[0008] FIG. 3 is a flow chart illustrating one example of a rule-based scoring method for a keyword evaluation system in which the content classes are phonemes;

[0009] FIG. 4 illustrates a method for providing feedback to the user of an automatic speech recognition system representing a quality of a keyword recorded by a user; and

[0010] FIG. 5 is a schematic block diagram illustrating an example system of hardware components capable of implementing examples of the systems and methods disclosed in FIGS. 1-4.

DETAILED DESCRIPTION OF EXAMPLE EMBODIMENTS

[0011] In many practical applications, a user may wish to train a speech recognizer device with a keyword of their choice. In this description, a “keyword” is a word or short phrase that is used by the user to request an action by the system via a speech recognition system. In many cases, a keyword is used to activate the system to receive a more complex command. User-recorded keywords are not always suitable for a given speech recognition application, resulting in excessive recognition errors and corresponding frustration for the user. The systems and methods described herein provide direct feed-back to the user about the quality of the enrolled command for keyword spotting purposes. This quality measure indicates to the user whether or not the enrolled keyword is likely to perform well. A good quality score reflects that the enrolled keyword is likely to be correctly recognized a high percentage of the time and that non-keyword

speech is not likely to be misrecognized as the keyword. A bad quality measure reflects the opposite. This score can be directly fed back to the user during a training period for the keyword, allowing the user to either re-train with a different keyword or re-train with the same keyword with a better enunciation.

[0012] FIG. 1 illustrates one example of a system 10 for evaluating the quality of a keyword for an automated speech recognition system 20. The automated speech recognition system can be implemented as machine readable instructions stored on a non-transitory computer readable medium and executed by an associated processor, as dedicated hardware, or as a mix of software and dedicated hardware. A microphone 12 records a keyword spoken by a user. The speech recognition system 20 comprises a decoder 22 and a quality evaluation component 24. The decoder 22 recognizes all or a portion of the keyword, and the quality evaluation component 24 assigns a score to the keyword representing a quality of the keyword.

[0013] In this description, the quality of the keyword, or a score representing the quality, is a numerical or categorical parameter representing either or both of a recognition rate of the keyword at a decoder of the automatic speech recognition system and a false positive rate of the keyword at the decoder of the automatic speech recognition system. In one implementation, a weighted linear combination of these two values is used as the quality score. Feedback representing the quality score can be provided to the user at a user interface 26. The user interface 26 can interact with an output device 30 associated with the automated speech recognition system 20, such as a speaker or display, to provide the feedback to the user. In one implementation, the feedback is provided substantially in real-time, such that the user may adjust the keyword or its enunciation during training of the keyword in response to the provided feedback.

[0014] FIG. 2 illustrates another example of a system 50 for evaluating the quality of a keyword for an automated speech recognition system. To evaluate the keyword quality, the system 50 first classifies each of the feature frames used in training a keyword model into different content classes, such as speech phonetic types, and then, based on the resulting sequence of content classes, assigns a quality score. Scoring rules for providing the quality score will depend on the type of decoder, the characteristics of the algorithm implemented in the decoder, the selected content classes, and the key word spotting application targeted. For example, in applications where it is desired to have low false alarm rates, the scoring weighting

may be more towards keywords that have more vowels and vowels of different kinds, because vowels are spectrally rich in information and generally contain higher acoustic signal energy.

[0015] The system 50 includes a microphone 52 that records a keyword spoken by a user, which is provided to a non-transitory computer readable medium 60 storing machine readable instructions, executable by an associated processor 54, to provide an automatic speech recognition system 60. A front end 62 of the speech recognition system divides the recorded keyword into a plurality of subunits. The recognizer front-end 62 also extracts a set of features from each of the plurality of subunits representing the audio content of each subunit. In the illustrated implementation, the recorded keyword is divided into frames of equal duration, such as ten milliseconds, to provide the subunits.

[0016] The recognizer front-end 62 converts the speech signal on a frame-by-frame basis to a set of feature vectors. The feature vectors represent a sufficient set of statistics for that frame of speech, and can include any of a variety of different representations, such as Mel-Frequency cepstral coefficients (MFCC), perceptual linear prediction coefficients (PLP), and linear predictive coding coefficients (LPC). In another implementation, one or more measures of the pitch, tone and energy of the recorded speech can be utilized as features.

[0017] The extracted sets of features for each frame are then provided to a decoder 64 that assigns one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit. In one example, a model of the chosen keyword is generated from feature vectors generated during training. Then during recognition, the feature vectors are input to the decoder (implemented, for example, as a hidden Markov model, a convolutional neural network, or a deep neural network) and compared with one or more models to recognize the spoken keyword. In accordance with an aspect of example embodiments, the content classes can either be parameters already provided by the decoder 64 during normal operation, or the models at the decoder can be modified to provide the content classes in addition to the normal detection function.

[0018] The specific content classes selected will vary with the application. In one implementation, the plurality of classes can simply represent respective quality scores including at least a first class, representing a first quality score, and a second class, representing a second quality score that is different from the first quality score. In such a case, each frame is simply assigned a score based on its extracted features and, in some applications, the features or

assigned classes of surrounding subunits. In another implementation, each of the content classes represents a phoneme being spoken by a user during the frame. In some systems, the detector 64 may be able to provide a phoneme for each frame as part of the existing recognition task. Specifically, in a system utilizing a model of the keyword and a filler or garbage model to identify or reject key words, the phonetic type can be determined using parameters generated in the filler model.

[0019] For the detection of phonemes, the set of phonetic types used may depend on the ability of the decoder 64 to reliably classify them as such. In one embodiment, the set of phonetic types can include onsets, fricatives, nasals, vowels, offsets, and silence. However, this set of phonemes is merely provided as an example, and other sets of phonemes and corresponding scoring systems can be utilized. A variety of methods are possible for mapping the feature frame, at a time n , to a corresponding phonetic type.

[0020] In a minimum distance based method, each phonetic type is first associated with a representative set of vectors. The idea is to find the vector across all the phonetic types that is closest to the input feature frame in terms of some distance measure. The length of the representative phonetic vector needs to be the same as that of the feature frame representation at the output of the front-end 62. The phonetic set of vectors can represent each phonetic type with different numbers of feature vectors and thus may be represented as:

$$\begin{aligned}\{\text{Onsets}\} &= [On_1, On_2, K, On_{k_1}] \\ \{\text{Fricatives}\} &= [Fr_1, Fr_2, K, Fr_{k_2}] \\ \{\text{Nasals}\} &= [Na_1, Na_2, K, Na_{k_3}] \\ \{\text{Vowels}\} &= [Vo_1, Vo_2, K, Vo_{k_4}] \\ \{\text{Offsets}\} &= [Off_1, Off_2, K, Off_{k_5}] \\ \{\text{Silence}\} &= [Si_1, Si_2, K, Si_{k_6}]\end{aligned}$$

[0021] In this representation, a total number of phonetic vectors is:

$$N = k_1 + k_2 + k_3 + k_4 + k_5 + k_6.$$

If $Ph_i \in \text{PhoneticSet}$ $i \in 1, \dots, N$ denotes a vector in the phonetic set, the minimum distance based rule classifies an input feature frame vector $FF(n)$ as one of onset, fricative, nasal, vowel, offset or silence based on the phonetic vector with the minimum p-norm error:

$$\text{PhoneticType}(n) = \arg \min_{i \in 1, \dots, N} (\|Ph_i - FF(n)\|_p) \quad \text{Eq. 1}$$

[0022] To simplify calculations, the 2-norm can be used, such that $p = 2$. In applications in which hidden Markov models are used, the hypothesis test is usually between a keyword model and a filler or garbage model. In this case, the vectors utilized in the filler model may correspond directly to the vectors in the phonetic set.

[0023] In the minimum distance method, the phonetic type at time instant n was chosen based on the feature frame of the input speech signal at time instant n . Alternatively, methods exist which take as input multiple feature frames (at different time instances) in order to make a decision on the phonetic type at time n , such that:

$$PhoneticType(n) = Function(FF(n - K), K, FF(n), K, FF(n + K)) \quad \text{Eq. 2}$$

[0024] These methods try to employ the context of previous and future frames in order to make a decision on the current frame. One such example is neural network implementation, which assigns likelihood probabilities to all the phonetic types. A simple hard decision rule at the output of the neural network is to choose the phonetic type at the output with the largest likelihood probability.

[0025] Each element of the phonetic set may not be a single vector. In some embodiments, each element of the set may contain parameters that describe a probability distribution of frame features that correspond to the phonetic element. Then, when one or more feature frames are observed, the phonetic set element parameters are used to calculate the probabilistic likelihood that the observed feature frames correspond to each of the elements of the phonetic set. A similar simple hard decision rule as described for neural networks can be used, which is to choose the phonetic type that provides the largest likelihood. Other more complex decision rules may also be applied by taking into account sets of higher likelihoods for one or more input feature frames.

[0026] A quality evaluation component 66 calculates a score representing a quality of the keyword from quality scores assigned to the plurality of subunits. In an embodiment in which each content class represents a quality score or can be directly associated with a quality score, the quality evaluation component can simply provide a weighted linear combination of the quality scores. For example, where the content classes represent phonemes, each phoneme can represent a specific value, with values across a syllable summed to provide the quality score for that syllable. In one embodiment, all scores can be equally weighted, although weights could be applied to the quality scores according to the position of the frame within the recorded word. In

other embodiments, the weights applied to each quality score can be derived from the content class of its associated frame and the content classes of surrounding frame.

[0027] In another embodiment, the quality evaluation component 66 utilizes a rule-based scoring system, in which each frame is assigned a score based on its associated content class as well as one or more other parameters. These parameters can be external to the content class determination, such as a tone, energy, or pitch of the frame, or can be drawn from the content classes of surrounding frames. In one embodiment, the quality evaluation component 66 assigns the quality score to each frame via a rule-based scoring in which the score for each frame is determined according to the phoneme associated with the frame as well as a phoneme associated with at least one adjacent frame.

[0028] FIG. 3 is a flow chart illustrating one example of a rule-based scoring method 100 for a keyword evaluation system in which the content classes are phonemes. In such a system, the mapping of the training model's feature vectors to the phonetic types results in a sequence of phonetic types. For a training model consisting of M feature vectors, the resultant phonetic vector model may be given as $phoneticModel = [Ph(1), K, Ph(M)]$. Each model is then broken into potentially multiple segments based on the phonetic class-type to which the input frame maps. Each syllabic segment is given a score, and the scoring system will vary with the desired application and the capabilities of the decoder 64. For example, in some embodiments, the scoring rules are biased toward putting more emphasis on sounds that have more vowels to reduce false positives in the keyword recognition.

[0029] At 102, a start of a segment is detected. The start of a segment may be indicated by an onset, a fricative, a nasal, and a vowel. At 104, a next feature vector is selected. At 106, the score, S , of the segment is initialized to zero. Other counters, such as the vowel counter described hereinbelow, can be initialized to their various start values at this point as well. At 108, it is determined if the selected feature vector is a vowel, fricative, nasal, or onset phoneme. If the feature vector is not any of these phonemes (OTHER), the method advances to 109. At 109, it is determined if the segment has ended. The end of a segment may be indicated by an offset, silence, or an end of the model, that is, the last feature vector. If the segment has ended (Y), the method terminates. Otherwise (N), the method returns to 104 to select a next feature vector.

[0030] If the selected feature vector is a vowel (VOWEL), it is then determined at 110 if the

identified vowel is a same vowel as a previous feature vector. If so (Y), the method advances to 112, where a counter, V, is incremented by one. If not (N), the method advances to 114, where the counter is reset to a value of one. Regardless, the method then proceeds to 116, where it is determined if the counter exceeds a threshold value. In this example, the threshold is four, but this threshold will vary with the desired application and the characteristics of the decoder 64. If the counter has a value higher than four (Y), the method advances to 109. If the counter has a value less than or equal to four (N), the score is incremented by a first value, X, at 118, and the method advances to 109. For example, X can be set to 0.1.

[0031] Returning to 108, if the selected feature vector is a fricative, nasal, or onset (NON-VOWEL), it is then determined at 120 if the identified phoneme is a same phoneme as a previous feature vector. This does not refer just to the same general type, but to the specific fricative, nasal, or onset, as represented by the various representative phonetic vectors used in the classification task. If so (Y), the method advances to 109. If the phoneme is distinct (N), the score is incremented by a second value, Z, at 122, and the method advances to 109. For example, the value Z can be set to 0.1. Once the method terminates, the score, S, represents a quality of the segment. Multiple segments can be added to obtain an overall quality for a keyword, such that a total quality score is then computed at the quality evaluation component 66 by adding up the individual syllabic segment scores and normalizing with respect to a constant, so that the scores lie in a range between 0 and 1.

[0032] The computed score can be provided to a feedback generation component 68 to be provided to the user at an output device, such as a speaker 70. In one implementation, the score classified into a categorical parameter, such as “good,” ”bad” or “average,” with this categorical parameter provided as feedback. Accordingly, the quality of the key word can be communicated to the user in a manner readily comprehensible without knowledge of the specific scoring system. In one implementation, each of the front end 62, the decoder 64, and the quality evaluation component 66 are designed such that the speaker 70 provides the feedback substantially in real time. This design allows the user to immediately record a new keyword when negative feedback is received.

[0033] In view of the structural and functional features described hereinabove, a methodology in accordance with various aspects of example embodiments is described with reference to FIG. 4. For simplicity of explanation, the methodology of FIG. 4 is shown and described as executing

serially, but example embodiments are not limited by the illustrated order, as some aspects could occur in different orders and/or concurrently with other aspects from that shown and described herein. Moreover, not all illustrated features may be required to implement a methodology in accordance with example embodiments.

[0034] FIG. 4 illustrates a method 150 for providing feedback to the user of an automatic speech recognition system representing a quality of a keyword recorded by a user. At 152, an input is received from a user at an automatic speech recognition system indicating that the user intends to submit a new keyword to the system. This procedure can be accomplished as an existing keyword command or via a software application associated with the recognition system. At 154, a keyword spoken by the user is recorded.

[0035] At 156, a score is assigned to the keyword representing a quality of the keyword, with the quality of the keyword representing either or both of a recognition rate of the keyword at a decoder of the automatic speech recognition system and a false positive rate of the keyword at the decoder of the automatic speech recognition system. In one embodiment, the recorded keyword into a plurality of subunits, and a set of features is extracted from each of the plurality of subunits. One of a plurality of content classes is assigned to each of the plurality of subunits according to at least the extracted set of features for each subunit, and the score representing the quality of the keyword from the content classes assigned to the plurality of subunits.

[0036] In one embodiment, a quality score is assigned to each of the plurality of subunits according to at least its assigned content class and the quality scores are combined across the plurality of subunits to provide the score representing the quality of the keyword. For example, each subunit can be associated with one of a plurality of phonemes, with each phoneme having an associated quality score, such that a first phoneme has a first quality score and a second phoneme has a second quality score different from the first quality score. Alternatively, each subunit can be associated with one of the plurality of phonemes, and a quality score is assigned to each subunit according to the phoneme associated with the subunit as well as a phoneme associated with at least one adjacent subunit. In one example, a first score is assigned to each subunit representing a vowel if less than a threshold number of consecutive preceding subunits represented the same vowel, and a second score is assigned to each subunit representing one of a fricative, a nasal, and an onset if the one of a fricative, a nasal, and an onset is different from the phoneme assigned in an immediately preceding subunit. A total score can be assigned from the

scores for the individual subunits, and substantially real-time feedback representing the assigned score can be provided to the user at 158.

[0037] FIG. 5 is a schematic block diagram illustrating an example system 200 of hardware components capable of implementing examples of the systems and methods disclosed in FIGS. 1-4. The system 200 can include various systems and subsystems. The system 200 can be a personal computer, a laptop computer, a workstation, a computer system, an appliance, an application-specific integrated circuit (ASIC), a field programmable gate array (FPGA), a server, a server blade center, a server farm, or other computing device.

[0038] The system 200 can include a system bus 202, a processing unit 204, a system memory 206, memory devices 208 and 210, a communication interface 212 (*e.g.*, a network interface), a communication link 214, a display 216 (*e.g.*, a video screen), and an input device 218 (*e.g.*, a keyboard and/or a mouse). The system bus 202 can be in communication with the processing unit 204 and the system memory 206. The additional memory devices 208 and 210, such as a hard disk drive, server, stand alone database, or other non-volatile memory, can also be in communication with the system bus 202. The system bus 202 interconnects the processing unit 204, the memory devices 206-210, the communication interface 212, the display 216, and the input device 218. In some examples, the system bus 202 also interconnects an additional port (not shown), such as a universal serial bus (USB) port.

[0039] The processing unit 204 can be a computing device and can include an application-specific integrated circuit (ASIC). The processing unit 204 executes a set of instructions to implement the operations of examples disclosed herein. The processing unit can include at least one processing core. The additional memory devices 206, 208 and 210 can store data, programs, instructions, database queries in text or compiled form, and any other information that can be needed to operate a computer. The memories 206, 208 and 210 can be implemented as computer-readable media (integrated or removable) such as a memory card, disk drive, compact disk (CD), or server accessible over a network. In some examples, the memories 206, 208 and 210 can comprise text, images, video, and/or audio, portions of which can be available in formats comprehensible to human beings. Additionally or alternatively, the system 200 can access an external data source or query source through the communication interface 212, which can communicate with the system bus 202 and the communication link 214.

[0040] In operation, the system 200 can be used to implement one or more parts of a keyword

evaluation system in accordance with example embodiments. Computer executable logic for evaluating a quality of the keyword resides on one or more of the system memory 206, and the memory devices 208, 210 in accordance with certain examples. The processing unit 204 executes one or more computer executable instructions originating from the system memory 206 and the memory devices 208 and 210. The term "computer readable medium" as used herein refers to a set of one or more non-transitory media that participate in providing instructions to the processing unit 204 for execution. These media can be local to the process or connected via a local network or Internet connection.

[0041] Modifications are possible in the described embodiments, and other embodiments are possible, within the scope of the claims.

CLAIMS

What is claimed is:

1. An automated speech recognition system, comprising:
 - a microphone that records a keyword spoken by a user;
 - a front end that divides the recorded keyword into a plurality of subunits, each containing a segment of recorded audio, and extracts a set of features from each of the plurality of subunits;
 - a decoder that assigns one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit; and
 - a quality evaluation component that calculates a score representing a quality of the keyword from the content classes assigned to the plurality of subunits.
2. The system of claim 1, further comprising an output device that provides feedback on the quality of the keyword to the user, each of the front end, the decoder, and the quality evaluation component being configured such that the output device provides the feedback substantially in real time.
3. The system of claim 1, wherein the plurality of features for each subunit comprise at least one of a pitch, an energy, a spectral representation of the audio, and a function of one of the pitch, the energy, and the spectral representation of the audio within the subunit.
4. The system of claim 1, wherein the decoder associates each subunit with one of a plurality of phonemes, with each phoneme having an associated quality score, and combines the quality scores across at least a subset of the plurality of subunits to provide the score representing the quality of the keyword.
5. The system of claim 1, wherein the front end divides the keyword into a plurality of frames, each frame having a same duration.
6. The system of claim 5, wherein the decoder associates each frame with one of a plurality of phonemes, with each phoneme having an associated quality score, and combines the quality scores across at least a subset of the plurality of frames to provide the score representing the quality of the keyword.
7. The system of claim 5, wherein the decoder associates each frame with one of a plurality of phonemes, the quality evaluation component assigning a quality score to each frame according to the phoneme associated with the frame as well as a phoneme associated with at least one adjacent frame.

8. The system of claim 7, wherein the quality evaluation component assigns a first score for each frame representing a vowel, and a second score for each frame representing one of a fricative, a nasal, and an onset.

9. The system of claim 8, wherein the quality evaluation component assigns the first score for a frame representing a given vowel only if less than a threshold number of consecutive preceding frames represented the given vowel, and the second score is assigned for a frame representing one of a fricative, a nasal, and an onset only if the one of a fricative, a nasal, and an onset is different from the phoneme assigned to an immediately previous frame.

10. A method comprising:

receiving an input from a user at an automatic speech recognition system indicating that the user intends to submit a new keyword to the system;

recording the keyword spoken by the user;

assigning a score to the keyword indicative of a quality of the keyword, the quality of the keyword representing at least one of a recognition rate of the keyword at a decoder of the automatic speech recognition system and a false positive rate of the keyword at the decoder of the automatic speech recognition system; and

providing substantially real-time feedback representing the assigned score to the user.

11. The method of claim 10, further comprising:

dividing the recorded keyword into a plurality of subunits;

extracting a set of features from each of the plurality of subunits; and

assigning one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit;

wherein assigning the score to the keyword representing the quality of the keyword comprises determining the score representing the quality of the keyword from the content classes assigned to the plurality of subunits.

12. The method of claim 11, wherein determining the score representing the quality of the keyword from the content classes assigned to the plurality of subunits comprises:

assigning a quality score to each of the plurality of subunits according to at least its assigned content class; and

combining the quality scores across the plurality of subunits to provide the score representing the quality of the keyword.

13. The method of claim 11, wherein assigning one of the plurality of content classes to each of the plurality of subunits comprises associating each subunit with one of a plurality of phonemes, with each phoneme having an associated quality score, such that a first phoneme has a first quality score and a second phoneme has a second quality score different from the first quality score.

14. The method of claim 11, wherein assigning one of the plurality of content classes to each of the plurality of subunits comprises associating each subunit with one of a plurality of phonemes and assigning a quality score to each subunit according to the phoneme associated with the subunit as well as a phoneme associated with at least one adjacent subunit.

15. The method of claim 14, wherein the quality evaluation component assigns the quality score for a given frame as a function of the phonemes assigned to at least one prior subunit and the phoneme assigned to the given subunit.

16. An automated speech recognition system, comprising:

- a microphone that records a keyword spoken by a user;
 - a front end that divides the recorded keyword into a plurality of subunits, each containing a segment of recorded audio, and extracts a set of features from each of the plurality of subunits;
 - a decoder that assigns one of a plurality of content classes to each of the plurality of subunits according to at least the extracted set of features for each subunit;
 - a quality evaluation component that calculates a score representing a quality of the keyword from the content classes assigned to the plurality of subunits; and
 - an output device that provides feedback on the quality of the keyword to the user;
- wherein each of the front end, the decoder, and the quality evaluation component are configured such that the output device provides the feedback substantially in real time.

17. The system of claim 16, wherein a first content class of the plurality of content classes represents a first quality score and a second content class of the plurality of content classes represents a second quality score that is different from the first quality score, the quality evaluation component determining the score representing a quality of the keyword by combining the quality scores assigned to each of the plurality of sub-units.

18. The system of claim 16, wherein the decoder associates each subunit with one of a plurality of phonemes, with each phoneme having an associated quality score, and combines the quality scores across at least a subset of the plurality of subunits to provide the score representing

the quality of the keyword.

19. The system of claim 18, wherein the quality evaluation component assigns a first score for each subunit representing a vowel, and a second score for each subunit representing a non-vowel phoneme.

20. The system of claim 16, wherein the decoder associates each subunit with one of a plurality of phonemes, the quality evaluation component assigning a quality score to each subunit according to the phoneme associated with the subunit as well as a phoneme associated with at least one adjacent subunit.

1/3




FIG. 1




FIG. 2

2/3




FIG. 3

3/3




FIG. 4




FIG. 5