



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) **EP 0 965 978 B9**

(12) **CORRECTED EUROPEAN PATENT SPECIFICATION**

Note: Bibliography reflects the latest situation

(15) Correction information:
Corrected version no 1 (W1 B1)
Corrections, see page(s) 17-19

(51) Int Cl.7: **G10L 15/06**

(48) Corrigendum issued on:
08.06.2005 Bulletin 2005/23

(45) Date of publication and mention
of the grant of the patent:
26.01.2005 Bulletin 2005/04

(21) Application number: **99304623.4**

(22) Date of filing: **15.06.1999**

(54) **Non-interactive enrollment in speech recognition**

Nicht-interaktive Registrierung zur Spracherkennung

Enregistrement non interactif pour reconnaissance de la parole

(84) Designated Contracting States:
DE FR GB

(30) Priority: **15.06.1998 US 94609**

(43) Date of publication of application:
22.12.1999 Bulletin 1999/51

(73) Proprietor: **Dragon Systems Inc.**
Newton, Massachusetts 02160 (US)

(72) Inventors:
• **Sherwood, Stefan**
Belmont, MA 02178 (US)
• **Parmenter, David Wilsberg**
Newton, MA 02158 (US)
• **Gould, Joel**
Winchester, MA 01890 (US)
• **Albina, Toffee A.**
Cambridge, MA 02139 (US)
• **Gold, Alan**
Acton, MA 01720 (US)

(74) Representative: **Howe, Steven**
Lloyd Wise
Commonwealth House,
1-19 New Oxford Street
London WC1A 1LW (GB)

(56) References cited:
EP-A- 0 867 857 **DE-A- 3 518 337**
US-A- 5 251 283 **US-A- 5 398 220**
US-A- 5 715 367

- **"Dragon NaturallySpeaking Professional", User guide, version 3.0, July 1998, Dragon Systems, Inc. XP002116949**
- **"EXAMPLE ENROLLMENT TEXT PLAYBACK FOR AN AUTOMATIC SPEECH RECOGNIZER" IBM TECHNICAL DISCLOSURE BULLETIN, vol. 36, no. 3, 1 March 1993 (1993-03-01), page 413 XP000354828 ISSN: 0018-8689**

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

EP 0 965 978 B9

Description

[0001] The invention relates to enrollment in speech recognition.

[0002] A speech recognition system analyzes a user's speech to determine what the user said. Most speech recognition systems are frame-based. In a frame-based system, a processor divides a signal descriptive of the speech to be recognized into a series of digital frames, each of which corresponds to a small time increment of the speech.

[0003] A continuous speech recognition system can recognize spoken words or phrases regardless of whether the user pauses between them. By contrast, a discrete speech recognition system recognizes discrete words or phrases and requires the user to pause briefly after each discrete word or phrase. Continuous speech recognition systems typically have a higher incidence of recognition errors in comparison to discrete recognition systems due to complexities of recognizing continuous speech.

[0004] In general, the processor of a continuous speech recognition system analyzes "utterances" of speech. An utterance includes a variable number of frames and may correspond to a period of speech followed by a pause of at least a predetermined duration.

[0005] The processor determines what the user said by finding acoustic models that best match the digital frames of an utterance, and identifying text that corresponds to those acoustic models. An acoustic model may correspond to a word, phrase or command from a vocabulary. An acoustic model also may represent a sound, or phoneme, that corresponds to a portion of a word. Collectively, the constituent phonemes for a word represent the phonetic spelling of the word. Acoustic models also may represent silence and various types of environmental noise.

[0006] The words or phrases corresponding to the best matching acoustic models are referred to as recognition candidates. The processor may produce a single recognition candidate (i.e. a single sequence of words or phrases) for an utterance, or may produce a list of recognition candidates.

[0007] An acoustic model generally includes data describing how a corresponding speech unit (e.g. a phoneme) is spoken by a variety of speakers. To increase the accuracy with which an acoustic model represents a particular user's speech, and thereby to decrease the incidence of recognition errors, the speech recognition system may modify the acoustic models to correspond to the particular user's speech. This modification may be based on samples of the user's speech obtained during an initial enrolment session and during use of the system.

[0008] Enrolment sessions for previous speech recognition systems typically required a user to read from a list of words or to read specific words in response to prompts. For example, DragonDictate® for Windows®, available from Dragon Systems, Inc. of Newton, Massachusetts, included a quick enrolment session that prompted a new user to speak each word of a small set of words, and then adapted the acoustic models based on the user's speech.

[0009] Other enrolment approaches also have been used. For example, NaturallySpeaking®, also available from Dragon Systems, Inc. of Newton, Massachusetts, includes an interactive enrolment session in which a new user recites a selected enrolment text. An associated display (e.g. an arrow) indicates the user's position in the text.

[0010] EP-A-0867857 is one of the applicant's earlier patent applications published after the priority date of the present application, and is therefore relevant for novelty considerations only. This document discloses a system for enrolment in a speech recognition system in which user utterances are matched to a portion of an enrolment text, and matching utterances are used to update acoustic models. The enrolment text may include variations, for example punctuation, which may or may not be uttered by a user. When matching a user utterance with the enrolment text, if the user utters such optional variations, these are used to update the acoustic model, however if the prescribed variations are not uttered, these are skipped.

[0011] US 5,251,283 relates to the training of a speech recognition system for use with a dictating machine. In particular, the disclosure relates to overcoming the problem of training the system using a microphone connected to the computer, but then using the system with speech recorded on a dictating machine, resulting in different acoustic properties of the speech to be recognised. The solution proposed is to use an interactive system in which the sound is recorded on a recording device in response to an interactive display or keyboard input, and is immediately transmitted to the speech recognition system. Therefore, the speech for training passes along the same path as the speech for subsequent recognition.

[0012] The invention, as defined in the appended dependent claims, provides non-interactive techniques for enrolling a user in a speech recognition system. Since the techniques are not interactive, the user may record enrolment speech using, for example, a portable recording device, and may later download the speech for processing to refine acoustic models of a speech recognition system. The techniques require the speech to generally correspond to an enrolment text, but permit the user to skip or repeat words, phrases, sentences, or paragraphs of the enrolment text. The techniques involve analyzing the user's speech relative to the enrolment text to identify portions of the speech that match portions of the enrolment text, and updating acoustic models corresponding to the matched portions of the enrolment text using the matching portions of the user's speech. The techniques promise to provide increased flexibility to the enrolment process, and to thereby simplify enrollment.

[0013] In one general aspect, a computer enrolls a user in a speech recognition system by obtaining data representing

speech by the user and generally corresponding to an enrollment text. The computer analyzes acoustic content of a user utterance, and, based on the analysis, determines whether the user utterance matches a portion of the enrollment text. If the user utterance matches a portion of the enrollment text, the computer uses the acoustic content of the user utterance to update acoustic models corresponding to the portion of the enrollment text. A determination that the user utterance matches a portion of the enrollment text is permitted even when the user has skipped or repeated words, sentences, or paragraphs of the enrollment text.

[0014] Implementations may include one or more of the following features. The enrollment procedure is not performed interactively. This means that the data representing the user's speech may be data recorded using a recording device physically separate from the computer. For example, the recording device may be a digital recording device, and obtaining data may include receiving a file from the digital recording device. Obtaining data also may include receiving signals generated by playing back the user's speech using a recording device, such as an analog recording device.

[0015] Prior to analyzing a user utterance, the computer may divide the data into groups, with each group representing an utterance by the user.

[0016] The computer may designate an active portion of the enrollment text, and may analyze acoustic content of an utterance relative to the active portion of the enrollment text. The computer may identify a position of a previously analyzed utterance in the enrollment text, and may designate the active portion of the enrollment text based on the identified position. The active portion may include text preceding and following the identified position. For example, the active portion may include a paragraph including the position, as well as paragraphs preceding and following that paragraph.

[0017] The computer may attempt to match the utterance to models for words included in the active portion of the enrollment text. To this end, the computer may employ an enrollment grammar corresponding to the active portion of the enrollment text.

[0018] A rejection grammar may be used to determine whether the user utterance matches a portion of the enrollment text. The rejection grammar may be a phoneme grammar and may model an utterance using a set of phonemes that is smaller than a set of phonemes used by the enrollment grammar.

[0019] The enrollment text may be selected from a group of enrollment texts, with each of the enrollment texts having a corresponding enrollment grammar. An enrollment text also may be received from a user. An enrollment grammar corresponding to the received enrollment text may be generated for use in determining whether the user utterance matches a portion of the enrollment text.

[0020] The user utterance may be ignored if it does not match a portion of the enrollment text.

[0021] In another general aspect, a user may be enrolled in a speech recognition system by recording the user's speech using a portable recording device and transferring the recorded speech to a computer. The computer then is used to analyze acoustic content of the recorded speech, identify, based on the analysis, portions of the speech that match portions of the enrollment text, and update acoustic models corresponding to matched portions of the enrollment text using acoustic content of matching portions of the speech. The recorded speech may skip or repeat portions of the enrollment text.

[0022] Other general aspects include obtaining data corresponding to enrollment text using a physically separate recording device, as well as designating an active portion of the enrollment text and analyzing acoustic content of an utterance relative to the active portion of the enrollment text.

[0023] A preferred example of a system according to the present invention will be described in accordance with the accompanying drawings, in which:

Fig. 1 is a block diagram of a speech recognition system.

Fig. 2 is a block diagram of speech recognition software of the system of Fig. 1.

Figs. 3 and 4 show a user interface for an enrollment program.

Figs. 5-7 are flow charts of enrollment procedures.

Fig. 8 is a flow chart of a model-adaptation procedure.

[0024] Referring to Fig. 1, a speech recognition system 100 includes input/output (I/O) devices (e.g., microphone 105, mouse 110, keyboard 115, and display 120) and a general purpose computer 125 having a processor 130, an I/O unit 135 and a sound card 140. A memory 145 stores data and programs such as an operating system 150, an application program 155 (e.g., a word processing program), and speech recognition software 160.

[0025] The microphone 105 receives the user's speech and conveys the speech, in the form of an analog signal, to the sound card 140, which in turn passes the signal through an analog-to-digital (A/D) converter to transform the analog signal into a set of digital samples. Under control of the operating system 150 and the speech recognition software 160, the processor 130 identifies utterances in the user's continuous speech. Utterances are separated from one another by a pause having a sufficiently-large, predetermined duration (e.g., 160-250 milliseconds). Each utterance may include one or more words of the user's speech.

[0026] The system also may include an analog recorder port 165 and/or a digital recorder port 170. The analog recorder port 165 is connected to the sound card 140 and is used to transmit speech recorded using a handheld recorder to the sound card. The analog recorder port may be implemented as a microphone positioned so as to be next to the speaker of the handheld recorder when the recorder is inserted into the port 165, and also may be implemented using the microphone 105, or as an analog cable connection. Alternatively, the analog recorder port 165 may be implemented as a tape player that receives a tape recorded using a handheld recorder and transmits information recorded on the tape to the sound card 140.

[0027] The digital recorder port 170 may be implemented to transfer a digital file generated using a handheld digital recorder. This file may be transferred directly into memory 145. The digital recorder port 170 may be implemented as a storage device (e.g., a floppy drive) of the computer 125, or as a modem or network card providing an electronic communication from a remote location.

[0028] Fig. 2 illustrates components of the speech recognition software 160. For ease of discussion, the following description indicates that the components carry out operations to achieve specified results. However, it should be understood that each component actually causes the processor 130 to operate in the specified manner.

[0029] Initially, a front end processing module 200 converts the digital samples 205 from the sound card 140 (or from the digital recorder port 170) into frames of parameters 210 that represent the frequency content of an utterance. Each frame includes 24 parameters and represents a short portion (e.g., 10 milliseconds) of the utterance.

[0030] Recognizer 215 receives and processes the frames of an utterance to identify text corresponding to the utterance. The recognizer entertains several hypotheses about the text and associates a score with each hypothesis. The score reflects the probability that a hypothesis corresponds to the user's speech. For ease of processing, scores are maintained as negative logarithmic values. Accordingly, a lower score indicates a better match (a high probability) while a higher score indicates a less likely match (a lower probability), with the likelihood of the match decreasing as the score increases. After processing the utterance, the recognizer provides the best-scoring hypotheses to the control/interface module 220 as a list of recognition candidates, where each recognition candidate corresponds to a hypothesis and has an associated score. Some recognition candidates may correspond to text while other recognition candidates correspond to commands.

[0031] The recognizer 215 processes the frames 210 of an utterance in view of one or more constraint grammars 225. A constraint grammar, also referred to as a template or restriction rule, may be a limitation on the words that may correspond to an utterance, a limitation on the order or grammatical form of the words, or both. For example, a constraint grammar for menu-manipulation commands may include only entries from the menu (e.g., "file", "edit") or command words for navigating through the menu (e.g., "up", "down", "top", "bottom"). Different constraint grammars may be active at different times. For example, a constraint grammar may be associated with a particular application program 155 and may be activated when the user opens the application program and deactivated when the user closes the application program. The recognizer 215 discards any hypothesis that does not comply with an active constraint grammar. In addition, the recognizer 215 may adjust the score of a hypothesis associated with a particular constraint grammar based on characteristics of the constraint grammar.

[0032] One constraint grammar 225 that may be used by the speech recognition software 160 is a large vocabulary dictation grammar. The large vocabulary dictation grammar identifies words included in the active vocabulary 230, which is the vocabulary of words known to the software. The large vocabulary dictation grammar also indicates the frequency with which words occur. A language model associated with the large vocabulary dictation grammar may be a unigram model that indicates the frequency with which a word occurs independently of context, or a bigram model that indicates the frequency with which a word occurs in the context of a preceding word. For example, a bigram model may indicate that a noun or adjective is more likely to follow the word "the" than is a verb or preposition.

[0033] The active vocabulary 230 uses a pronunciation model in which each word is represented by a series of phonemes that comprise the phonetic spelling of the word. Each phoneme may be represented as a triphone, a context-dependent phoneme, that includes three nodes. For example, the triphone "abc" represents the phoneme "b" in the context of the phonemes "a" and "c", with the phoneme "b" being preceded by the phoneme "a" and followed by the phoneme "c".

[0034] The phonemes are represented by acoustic models 235. In the case of triphones, the acoustic models 235 represent each triphone node as a mixture of Gaussian probability density functions ("PDFs"). For example, node "i" of a triphone "abc" may be represented as $ab^i c$:

$$ab^i c = \sum_k w_k N(\mu_k, c_k) ,$$

where each w_k is a mixture weight,

$$\sum_k w_k = 1,$$

5

μ_k is a mean vector for the probability density function ("PDF") N_k , and c_k is the covariance matrix for the PDF N_k . Like the frames in the sequence of frames, the vectors μ_k each include twenty four parameters. The matrices c_k are twenty four by twenty four matrices. Each triphone node may be represented as a mixture of up to, for example, sixteen different PDFs.

10 **[0035]** A complete dictation vocabulary consists of the active vocabulary 230 plus a backup vocabulary 245. The backup vocabulary may include files that contain user-specific backup vocabulary words and system-wide backup vocabulary words. The system-wide backup vocabulary contains all the words known to the system, including words which may currently be in an active vocabulary.

15 **[0036]** During error correction, word searches of the backup vocabularies start with the user-specific backup vocabulary and then check the system-wide backup vocabulary. The backup vocabularies also are searched when there are new words in text that a user has typed.

[0037] The recognizer 215 may operate in parallel with a pre-filtering procedure 240. Upon initiating processing of an utterance, the recognizer 215 requests from the pre-filtering procedure 240 a list of words that may have been spoken as the first word of the utterance (i.e., words that may correspond to the first and subsequent frames of the utterance). The pre-filtering procedure 240 performs a coarse comparison of the sequence of frames with the active vocabulary 230 to identify a subset of the vocabulary for which a more extensive comparison using the recognizer is justified.

25 **[0038]** The control/interface module 220 controls operation of the speech recognition software and provides an interface to other software or to the user. The control/interface module receives the list of recognition candidates for each utterance from the recognizer. Recognition candidates may correspond to dictated text, speech recognition commands, or external commands. When the best-scoring recognition candidate corresponds to dictated text, the control/interface module provides the text to an active application, such as a word processor. The control/interface module also may display the best-scoring recognition candidate to the user through a graphical user interface. The control/interface module controls operation of the speech recognition software in response to speech recognition commands (e.g., "wake up", "make that"), and forwards external commands to the appropriate software.

30 **[0039]** The control/interface module also controls the active vocabulary, acoustic models, and constraint grammars that are used by the recognizer. For example, when the speech recognition software is being used in conjunction with a particular application (e.g., Microsoft Word), the control/interface module updates the active vocabulary to include command words associated with that application and activates constraint grammars associated with the application.

35 **[0040]** Other functions provided by the control/interface module 220 include a vocabulary customizer and a vocabulary manager. The vocabulary customizer optimizes the language model of a specific topic by scanning user supplied text. The vocabulary manager is a developer tool which is used to browse and manipulate vocabularies, grammars and macros. Each such function of the control/interface module 220 may be implemented as an executable program that is separate from the main speech recognition software. Similarly, the control/interface module 220 also may be implemented as a separate executable program.

40 **[0041]** The control/interface module 220 also may provide an enrollment program that uses an enrollment text and a corresponding enrollment grammar to customize the speech recognition software to a specific user. The enrollment program may operate in an interactive mode that guides the user through the enrollment process, or in a non-interactive mode that permits the user to enroll independently of the computer. In the interactive mode, the enrollment program displays the enrollment text to the user and the user reads the displayed text. As the user reads, the recognizer 215 uses the enrollment grammar to match a sequence of utterances by the user to sequential portions of the enrollment text. When the recognizer 215 is unsuccessful, the enrollment program prompts the user to repeat certain passages of the text. The recognizer uses acoustic information from the user's utterances to train or adapt acoustic models 235 corresponding to the matched portions of the enrollment text.

50 **[0042]** In the non-interactive mode, the user reads the text without prompting from the computer. This offers the considerable advantage that, in addition to reading text displayed by the computer, the user can read from a printed text independently of the computer. Thus, the user could read the enrollment text into a portable recording device and later download the recorded information into the computer for processing by the recognizer, or could read the enrollment text over a telephone line. In addition, the user is not required to read every word of the enrollment text, and may skip words or paragraphs as desired. The user also may repeat portions of the text. This adds substantial flexibility to the enrollment process.

55 **[0043]** The enrollment program may provide a list of enrollment texts, each of which has a corresponding enrollment grammar, for the user's selection. Alternatively, the user may input an enrollment text from another source. In this case,

the enrollment program may generate the enrollment grammar from the input enrollment text, or may employ a previously generated enrollment grammar.

[0044] Fig. 3 shows a user interface 300 of the interactive enrollment program. The user interface 300 displays an enrollment text 310 for the user to read and displays the user's current position, also referred to as the "enrollment position", using, for example, a moving arrow 320 in the enrollment text. The display of Fig. 3 shows an enrollment position at the word "for", indicating that the enrollment program has recognized the user's reading of the enrollment text up to the word "for" and that the enrollment process will continue when the user continues reading, beginning with the word "for". By contrast, the display of Fig. 4 indicates an enrollment position at the word "program". In addition to using the arrow, the user interface may display the enrollment position by highlighting portions of the text as they are recognized or displaying a cursor at the enrollment position.

[0045] Using the user interface 300 of Fig. 3, the user starts the enrollment program by selecting the Record button 330 through a voiced command or an input device such as a mouse. The user then reads the displayed text 310 beginning with the text at the enrollment position, as indicated by the arrow 320. As the user reads, the recognizer 215 matches the user's speech to the enrollment text using the enrollment grammar and advances the enrollment position (and arrow 320) to the beginning of the next word to be spoken. Acoustic models corresponding to the enrollment text are updated based on the matching user utterances. In general, the recognizer 215 does not use the pre-filtering procedure 240 during the enrollment process.

[0046] As shown in Fig. 5, the enrollment program begins by setting the enrollment position at the beginning of the enrollment text (step 1510) and displaying the arrow at the enrollment position (step 1520). The enrollment program next receives the user's digitized speech for an utterance (step 1530). The recognizer 215, using the enrollment grammar corresponding to the enrollment text, determines whether the utterance matches a portion of the enrollment text beginning at the enrollment position (step 1540). Because the portion of the text to which an utterance corresponds is unspecified, the recognizer 215 uses the enrollment grammar to identify a portion of the enrollment text beginning at the enrollment position that best matches the utterance.

[0047] If there is no match, the enrollment program repeats the process of retrieving and comparing the user's speech with the enrollment text (steps 1530, 1540). If there is a match, the enrollment program updates the acoustic models 235 corresponding to the matched portion of the enrollment text based on the user utterance (step 1550). Alternatively, the enrollment program may update the acoustic models after specific segments (such as sentences or paragraphs) of enrollment text have been recognized, or after the entire enrollment text has been recognized.

[0048] If no enrollment text remains after the matched portion of enrollment text (step 1560), the enrollment program for that enrollment text ends (step 1570). If the enrollment program provides multiple enrollment texts, after finishing one enrollment text, the enrollment program may automatically begin the next enrollment text by displaying the new enrollment text and setting the enrollment position at the beginning of that new text (step 1510).

[0049] If enrollment text remains at step 1560, the enrollment position is advanced to the word following the matched portion of the enrollment text (step 1580), the arrow display is updated to point to the updated enrollment position (step 1520), and the enrollment process continues.

[0050] Different methods may be used to determine whether a user utterance matches a portion of the enrollment text (step 1540). For example, the enrollment grammar may be implemented to recognize that different users may pause at different positions in the enrollment text, with the result that utterances of different users may not uniformly correspond to the same portion of the enrollment text. As noted above, the enrollment grammar may handle such situations by allowing different lengths of enrollment text to match a user's utterance.

[0051] The enrollment grammar may be implemented to recognize and appropriately handle expected variations in the user's reading of the enrollment text in determining whether a user utterance matches a portion of the enrollment text. An example of such a variation is the dictation convention of enunciating punctuation marks: at the end of sentence, the user may enunciate the word "period". To handle this situation, the enrollment grammar may be implemented to recognize the word "period" at the end of sentences as optional. Thus, if the user says "period", the enrollment grammar matches the user's speech to the word period for purposes of updating the acoustic models. If the user does not say period, the enrollment grammar skips the optional word "period" and proceeds to the next word in the enrollment text.

[0052] Similar rules may be implemented for handling user enunciation of spacing (e.g., new line or new paragraph), and formatting conventions (e.g., capitalization, bold, underline, or italics). For example, the enrollment grammar for the passage of text illustrated in Fig. 3 may be expressed in Extendable Backus-Naur Form (EBNF) as:

```
<recognition result> ::= [Capitalize] dragon [Capitalize] systems is pleased to acknowledge [Capitalize] arthur [Capitalize] c [Period] [Capitalize] clarke and [Capitalize] harper [Capitalize] collins [Capitalize] publishers for allowing us to use selections from [Capitalize] arthur [Capitalize] c [Period] [Capitalize] {clarke's | clarke Apostrophe s} novel three thousand [and] one [Colon] [Capitalize] the [Capitalize] final [Capitalize] odyssey [Comma] for our training program [Period],
```

where

[] means optional,

< > means a rule,

! means an OR function, and

::= means "is defined as" or "is".

The enrollment grammar is modified as the user advances through the text. For example, when the user is at the word "for" in the text, as indicated by the arrow 320 in Fig. 3, the enrollment grammar may be expressed in EBNF as:

<recognition result> ::= for our training program [Period]

Alternatively, the enrollment grammar may be expressed as a set of rules. Appendix A includes a possible rule-based grammar for the enrollment text shown in Fig. 3. An advantage of the ruled-based grammar is that it does not need to be modified as the user advances through the text.

[0053] In general, use of the enrollment program reduces the recognition error rate of the speech recognition system by customizing acoustic models to a user's speech based on the user's pronunciation of sounds corresponding to those acoustic models in reciting the enrollment text. However, if the user says something other than the expected enrollment text, the user's speech should not be used to update the acoustic models corresponding to the expected enrollment text. This situation may arise, for example, if the user skips a line in the displayed text 310, mistakenly reads a displayed word as a different word, or is interrupted while reading the displayed text and engages in an outside conversation. This situation also may arise if the speech recognition system picks up sounds other than the user's speech, such as a cough or sneeze, or environmental noise such as traffic or a barking dog.

[0054] The enrollment grammar permits the recognizer 215 to find the portion of the enrollment text that best corresponds to the user's utterance. However, the enrollment grammar does not permit the recognizer 215 to verify that the utterance actually corresponds to the enrollment text. For this purpose, the enrollment program uses a rejection grammar in conjunction with the enrollment grammar to identify situations in which the sound detected by the recognizer does not match the expected enrollment text.

[0055] Fig. 6 elaborates the process of step 1540 of Fig. 5, where both the enrollment grammar and the rejection grammar are used to detect situations in which the user utterance should not be used to update the acoustic models for enrollment text. The enrollment program identifies a best recognition candidate that corresponds to the enrollment grammar and designates the score provided by the recognizer 215 for that candidate as the enrollment score (step 1541). The recognizer 215 may determine the score for the candidate based on criteria, such as matching acoustic models and word context.

[0056] Even the best recognition candidate may not be a good match to the enrollment text. For example, if the enrollment position is at "for", as shown in Fig. 3, and the user utterance is "from Arthur C. Clarke's novel" from the previous line, the recognizer 215 will attempt to match the utterance to the portion of the enrollment grammar starting with "for", and will produce a recognition candidate corresponding to the best match to that portion of the enrollment grammar.

[0057] The enrollment program also identifies a best recognition candidate that corresponds to the rejection grammar and designates a score provided by the recognizer 215 for that candidate as a rejection score (step 1542). In general, the rejection score should be better than the enrollment score when the user utterance does not correspond to the expected portion of the enrollment text.

[0058] The rejection grammar may be expressed in EBNF as:

<recognition result> ::= <phonemes>

where

<phonemes> ::= phoneme [<phonemes>]

and phoneme is any phoneme in a rejection grammar phoneme set. The rejection grammar phoneme set is selected such that most sounds in the expected speech can be reasonably modeled. If, as in the above example, the rejection grammar allows any combination of these phonemes, the processing required to find the best matching phonemes for an utterance can be expensive. Accordingly, the rejection grammar phoneme set may include only a subset of the set of phonemes used by the speech recognizer. For example, the attached Appendix B illustrates a set of fifty phonemes, with a subset of nineteen of the phonemes being included in the rejection grammar phoneme set.

[0059] The rejection grammar matches the acoustics of the user utterance to a sequence of phonemes from the rejection grammar phoneme set. The score associated with a rejection grammar recognition candidate reflects the degree to which the user utterance corresponds to the matched sequence of phonemes and is determined so that the rejection score is likely to be better than the enrollment score when the user utterance does not match a portion of the enrollment text.

[0060] Again referring to the above example of Fig. 3, in which the enrollment position is at "for" and the user utterance is "from Arthur C. Clarke's novel", the enrollment grammar will likely result in a recognition candidate with a score

reflecting a poor match. On the other hand, the rejection grammar will result in a recognition candidate that may not be "from Arthur C. Clarke's novel" exactly, but will likely be closer than the enrollment recognition candidate beginning with "for". Accordingly, the rejection grammar should result in a recognition candidate having a better match and a better score.

[0061] The enrollment program determines whether the user utterance matches an enrollment utterance by comparing the enrollment score to the rejection score (step 1543). The precise method of comparison may vary. For example, in a simple approach, the enrollment program compares the enrollment and rejection scores and, if the enrollment score is better than the rejection score, determines that the user utterance matches a portion of the enrollment text (step 1544). Alternatively, if the rejection score is better than the enrollment score, as in the above example, the enrollment program determines that the user utterance does not match a portion of the enrollment text (step 1545).

[0062] The rejection grammar has been described with reference to its use with the enrollment grammar. However, it also may be used in other circumstances, and, in general, is used whenever the large vocabulary dictation constraint grammar is not active. For example, the speech recognition software may include a "Go To Sleep" command that deactivates all constraint grammars except the rejection grammar and a grammar for recognizing a "Wake Up" command. In general, after the "Go To Sleep" command has been implemented, the recognizer 215 will recognize only the "Wake Up" command. The rejection grammar enables the control/interface module 220 to distinguish between the "Wake Up" command and noise or other sounds.

[0063] Referring to Fig. 7, the non-interactive enrollment program operates according to a procedure 1700. Initially, the program obtains user speech corresponding to an enrollment text (step 1705). This speech corresponds to the user's entire recitation of the enrollment text and may be obtained in a number of ways. For example, the user may record the speech by reading from a printed text using an analog recording device. The user then may provide the recorded speech to the computer 125 using the port 165. The printed text may be provided by the enrollment program.

[0064] The user also may record the speech using a digital recording device. In this case, the user may provide a digital file corresponding to the recorded speech using the port 170.

[0065] The user also may provide the speech using the microphone 105 and reading from a printed text or text displayed on the display 120. In contrast to the interactive enrollment program, the non-interactive enrollment program does not prompt the user as to the user's position in the text. Rather, the non-interactive enrollment program uses a text editor or comparable program to display the entire enrollment text and permits the user to navigate through the enrollment text as desired. Alternatively, the enrollment program could provide a printed version of the enrollment text for the user to read.

[0066] After the user speaks the enrollment text, the program processes the speech to produce sets of frames corresponding to user utterances (step 1710). When the speech is provided using the port 165 or the microphone 105, the program processes the speech using the sound card 140 and the procedure 300 as discussed above. When the speech is provided using the port 170, processing typically involves reformatting the contents of the digital file. In either case, the program may determine that the speech has ended when the program encounters silence of sufficient duration (e.g., thirty seconds).

[0067] Next, the enrollment program designates an active portion of the enrollment text and builds an enrollment constraint grammar for the active portion (step 1715). For example, the active portion of the enrollment text may initially be the first three to five paragraphs of the enrollment text. In this case, the enrollment program builds an enrollment grammar that includes constraint grammars for each of the designated number of paragraphs. In general, the constraint grammars are flexible and only require that the user has said a sequence of words from the enrollment text in order. For example, an utterance may begin with any word in the enrollment grammar, and may end with any word in the enrollment grammar, so long as any intermediate words of the utterance are found, in order, in the enrollment grammar between the starting and ending words. Thus, the enrollment grammar permits recognition to occur even when the user has skipped or repeated words, sentences, or even entire paragraphs.

[0068] The enrollment grammar for the first five paragraphs of an enrollment text may be expressed in EBNF as:

$$\langle \text{recognition result} \rangle ::= \langle \text{paragraph 1} \rangle \mid \langle \text{paragraph 2} \rangle \mid \langle \text{paragraph 3} \rangle \mid \langle \text{paragraph 4} \rangle \mid \langle \text{paragraph 5} \rangle,$$

where

$$\langle \text{paragraph } n \rangle ::= [\text{word}^{1n}[\text{word}^{2n}[\text{word}^{3n} \dots \text{word}^{mn}]]] \mid [\text{word}^{2n}[\text{word}^{3n} \dots \text{word}^{mn}]] \mid \dots [\text{word}^{mn}],$$

and

"wordⁱ_n" is word *i* from paragraph *n*.

As discussed above, the enrollment grammar may be implemented to recognize and appropriately handle expected variations in the user's reading of the enrollment text in determining whether a user utterance matches a portion of the enrollment text. Thus, as with the interactive enrollment grammar, the enrollment grammar accounts, for example, for punctuation marks, spacing, and formatting.

[0069] Next, the enrollment program retrieves an utterance (step 1720) and causes the recognizer to perform rec-

ognition on the utterance using the active enrollment grammar (step 1725). This recognition includes use of a rejection grammar and proceeds as discussed above with respect to the interactive enrollment program.

[0070] If the recognizer finds words in the utterance that match the enrollment text (step 1730), the enrollment program updates the acoustic models for the phonemes of those words (step 1735). Alternatively, the enrollment program may defer updating the acoustic models until all of the utterances have been analyzed.

[0071] After updating the acoustic models, the enrollment program determines whether the active portion of the enrollment text needs to be updated (step 1740). Generally, the enrollment program updates the active portion to maintain a window of active paragraphs around the current position of the user's speech in the enrollment text. Thus, for example, if the recognizer has just identified text that appears in the fourth paragraph of the enrollment text, and the enrollment program maintains a five paragraph active portion currently including the first through fifth paragraphs, the enrollment program may update the active portion to include paragraphs two through six.

[0072] If the active portion of the enrollment text needs to be updated, the enrollment program updates the active portion and the active enrollment grammar (step 1745). Thus, in the example noted above, the enrollment program would eliminate the portion of the constraint grammar corresponding to the first paragraph of the enrollment text and would add a portion corresponding to the sixth paragraph. The modified enrollment grammar may be expressed in EBNF as:

$$\langle \text{recognition result} \rangle ::= \langle \text{paragraph 2} \rangle | \langle \text{paragraph 3} \rangle | \langle \text{paragraph 4} \rangle | \langle \text{paragraph 5} \rangle | \langle \text{paragraph 6} \rangle.$$

[0073] If the processed utterance is not the last utterance (step 1750), the enrollment program retrieves the next utterance (step 1720) and repeats the process.

[0074] Upon processing the last utterance, the enrollment program displays a report of the enrollment performance (step 1755). The report may include an indication of the percentage of words in the user's speech that were recognized. The report also may include an indication of the expected recognition performance using the modified acoustic models. The user may use the information from this report to determine whether it would be productive to repeat the enrollment process or to enroll using a different text.

[0075] The enrollment program creates a speaker-adapted model for the user by adapting speaker-independent acoustic models corresponding to the enrollment text based on the user's utterances. As described above, the speaker-independent acoustic models represent each triphone node as a mixture of Gaussian probability density functions ("PDFs") N_k having mean vectors μ_k and covariance matrices c_k , with the contribution made to the triphone node by each PDF being represented by a mixture weight w_k . The speaker-adapted model uses the mixture weights w_k and the covariance matrices c_k of the speaker-independent model, but uses adapted mean vectors μ_{kA} .

[0076] Referring to Fig. 8, the enrollment program generates the adapted mean vectors μ_{kA} according to a procedure 1800. Starting with the mean vectors μ_k of the speaker-independent model (step 1805), the enrollment program uses a transformation-based approach to produce the adapted mean vectors μ_{kA} . The transformation-based approach assumes that an adapted mean vector μ_{kA} may be expressed as:

$$\mu_{kA} = A\mu_k + b,$$

where **A** and **b** are transforms. When each mean vector has 24 entries, **A** is a 24 x 24 matrix and **b** has 24 entries.

[0077] Next, the enrollment program generates collections of PDFs that are expected to share a common transformation (step 1810). A collection C could include the PDFs representing similarly-sounding triphones, such as all PDFs associated with vowels or all PDFs associated with consonants. Another approach to generating collections is to group PDFs having mean vectors μ_k with similar values. This approach permits simplified generation of large numbers of collections. In general, implementations may include from one to several hundred different collections. Initial experimental results indicate that the first approach provides better results (i.e. results with lower recognition error rates). In general, the collections may be identified prior to the enrollment session.

[0078] Next, for each collection, the enrollment program assembles all speaker data for the PDFs included in the collection (step 1815). This speaker data corresponds to the user utterances produced during the enrollment session.

[0079] Next, the enrollment program determines $y[j]$, the average value of the speaker data vector, for each PDF j (step 1820). If $f[n]$ is the vector for the n -th frame of speaker data, $p_j(n)$ is the probability that the n -th frame of speaker data corresponds to a PDF j (i.e., the mixture weight w_j for the triphone node associated with the n -th frame of speaker data), and N is the total number of frames of speaker data, then $y[j]$ and the frame count, $N[j]$, for the frame j may be determined as:

$$y[j] = \left(\sum_{n=1}^N p_j(n) \right)^{-1} \sum_{n=1}^N p_j(n) f[n],$$

$$N[j] = \sum_{n=1}^N p_j(n).$$

This is usually referred to as the Baum-Welch or EM algorithm.

[0080] Next, the enrollment program produces transforms ($\mathbf{A}_C$ and $\mathbf{b}_C$) for each collection C using the relationship between the average values of the speaker data and the PDFs from the speaker-independent model (step 1825). This relationship may be expressed generally as:

$$y[j] = \mathbf{A}_C x[j] + \mathbf{b}_C + e,$$

where $x[j]$ corresponds to the mean vector μ_j for a PDF j and e is an error term. Based on this relationship, $\mathbf{A}_C^T$, the transpose of the transform $\mathbf{A}_C$, may be determined as:

$$\mathbf{A}_C^T = \left(\sum_{j=1}^{N_C} N(j) x[j] x^T[j] \right)^{-1} \sum_{j=1}^{N_C} N(j) x[j] y^T[j],$$

where $x^T[j]$ is the transpose of $x[j]$, $y^T[j]$ is the transpose of $y[j]$, and N_C is the number of components in the collection C.

[0081] The transform $\mathbf{A}_C$ is then produced by transposing $\mathbf{A}_C^T$, and the transform $\mathbf{b}_C$ is determined as:

$$\mathbf{b}_C = \left(\sum_{j=1}^{N_C} N(j) \right)^{-1} \sum_{j=1}^{N_C} N(j) (y[j] - \mathbf{A}_C x[j]).$$

[0082] Using the transforms, the enrollment program determines adapted mean vectors, μ_{jA} , for each component in the collection C (step 1830). In particular, the adapted mean vectors are determined as:

$$\mu_{jA} = \mathbf{A}_C \mu_j + \mathbf{b}_C.$$

The enrollment program then stores the adapted mean vectors μ_{jA} along with a speaker-identifier as a speaker-adapted model for later use in recognizing speech by the user (step 1835). This process may be iterated several times by using the adapted models to update the frame counts, speaker ($N[j]$ and $y[j]$), and the transform matrix.

[0083] In another implementation, fractions of the speaker independent means ($x[j]$) and the speaker-specific means may be used to compute $y[j]$ and $N[j]$: where r controls the relative importance of the speaker-independent means and the observed data. The value of r is optimized as a function of the amount of the

$$y[j] = \left(r + \sum_{n=1}^N p_j(n) \right)^{-1} \left(\sum_{n=1}^N p_j(n) f[n] + r x[j] \right)$$

$$N[j] = \tau + \sum_{n=1}^N P_j(n)$$

adaptation data available.

[0084] Other embodiments are within the scope of the following claims. For example, the specific methods by which the enrollment and rejection grammars score a user utterance may vary. Similarly, the method by which the scores are evaluated to determine whether the user utterance matches a portion of the enrollment text may vary.

[0085] The techniques described here are not limited to any particular hardware or software configuration; they may find applicability in any computing or processing environment that may be used for speech recognition. The techniques may be implemented in hardware or software, or a combination of the two. Preferably, the techniques are implemented in computer programs executing on programmable computers that each include a processor, a storage medium readable by the processor (including volatile and non-volatile memory and/or storage elements), at least one input device, and one or more output devices. Program code is applied to data entered using the input device to perform the functions described and to generate output information. The output information is applied to one or more output devices.

[0086] Each program is preferably implemented in a high level procedural or object oriented programming language to communicate with a computer system. However, the programs can be implemented in assembly or machine language, if desired. In any case, the language may be a compiled or interpreted language.

[0087] Each such computer program is preferably stored on a storage medium or device (e.g., CD-ROM, hard disk or magnetic diskette) that is readable by a general or special purpose programmable computer for configuring and operating the computer when the storage medium or device is read by the computer to perform the procedures described in this document. The system may also be considered to be implemented as a computer-readable storage medium, configured with a computer program, where the storage medium so configured causes a computer to operate in a specific and predefined manner.

Appendix A

[0088]

5

10

15

20

25

30

35

40

45

50

55

Follows is an enrollment grammar for the paragraph used for examples in the patent document (the rule that is initially activated in this example is AltList0):

```

5 rule0=Dragon ( <rule1> )
  rule1=systems ( <AftAlt1> )
  rule2=dragon ( <rule3> )
  rule3=systems ( <AftAlt1> )
10 rule4=Dragon Systems ( <AftAlt1> )
  AltList0=<rule0> | <rule2> | <rule4>
  AftAlt1=is ( <rule7> )
  rule7=please ( <rule8> )
  rule8=to ( <rule9> )
  rule9=acknowledge ( <rule10> )
15 rule10=Arthur ( <AltList2> )
  rule11=C. ( <AftAlt3> )
  rule12=( C\carbon ) ( <rule13> )
  rule13=( .\dot ) ( <AftAlt3> )
  rule14=( C\carbon ) ( <rule15> )
  rule15=( .\period ) ( <AftAlt3> )
20 rule16=( C\carbon ) ( <rule17> )
  rule17=( .\point ) ( <AftAlt3> )
  rule18=( c\charlie ) ( <rule19> )
  rule19=( .\dot ) ( <AftAlt3> )
  rule20=( c\charlie ) ( <rule21> )
  rule21=( .\period ) ( <AftAlt3> )
25 rule22=( c\charlie ) ( <rule23> )
  rule23=( .\point ) ( <AftAlt3> )
  rule24=( C\Roman one hundred ) ( <rule25> )
  rule25=( .\dot ) ( <AftAlt3> )
  rule26=( C\Roman one hundred ) ( <rule27> )
  rule27=( .\period ) ( <AftAlt3> )
30 rule28=( C\Roman one hundred ) ( <rule29> )
  rule29=( .\point ) ( <AftAlt3> )
  AltList2=<rule11> | <rule12> | <rule14> | <rule16> | <rule18> | <rule20>
  | <rule22> | <rule24> | <rule26> | <rule28>
  AftAlt3=Clarke ( <rule32> )
35 rule32=and ( <AltList4> )
  rule33=harper ( <AftAlt5> )
  rule34=Harper ( <AftAlt5> )
  AltList4=<rule33> | <rule34>
  AftAlt5=Collins ( <rule37> )
  rule37=publishers ( <rule38> )
40 rule38=for ( <rule39> )
  rule39=allowing ( <rule40> )
  rule40=us ( <rule41> )
  rule41=to ( <rule42> )
  rule42=use ( <rule43> )
  rule43=selections ( <rule44> )
45 rule44=from ( <rule45> )
  rule45=Arthur ( <AltList6> )
  rule46=C. ( <AftAlt7> )
  rule47=( C\carbon ) ( <rule48> )
  rule48=( .\dot ) ( <AftAlt7> )
  rule49=( C\carbon ) ( <rule50> )
50 rule50=( .\period ) ( <AftAlt7> )
  rule51=( C\carbon ) ( <rule52> )
  rule52=( .\point ) ( <AftAlt7> )
  rule53=( c\charlie ) ( <rule54> )
  rule54=( .\dot ) ( <AftAlt7> )
  rule55=( c\charlie ) ( <rule56> )
55

```

```

rule56={ .\period } [ <AftAlt7> ]
rule57={ c\charlie } [ <rule58> ]
rule58={ .\point } [ <AftAlt7> ]
5 rule59={ C\Roman one hundred } [ <rule60> ]
rule60={ .\dot } [ <AftAlt7> ]
rule61={ C\Roman one hundred } [ <rule62> ]
rule62={ .\period } [ <AftAlt7> ]
rule63={ C\Roman one hundred } [ <rule64> ]
10 rule64={ .\point } [ <AftAlt7> ]
AltList6=<rule46> | <rule47> | <rule49> | <rule51> ; <rule53> : <rule55>
| <rule57> | <rule59> | <rule61> | <rule63>
AftAlt7=Clarke [ <AltList8> ]
rule67={ 's\apostrophe-ess } [ <AftAlt9> ]
rule68={ 's\iz\r } [ <AftAlt9> ]
15 rule69={ 's\s\r } [ <AftAlt9> ]
rule70={ 's\z\r } [ <AftAlt9> ]
rule71={ '\apostrophe } [ <rule72> ]
rule72={ s\sierra } [ <AftAlt9> ]
rule73={ '\close-single-quote } [ <rule74> ]
rule74={ s\sierra } [ <AftAlt9> ]
20 rule75={ '\open-single-quote } [ <rule76> ]
rule76={ s\sierra } [ <AftAlt9> ]
AltList8=<rule67> | <AftAlt9> | <rule68> | <rule69> | <rule70> |
<rule71> | <rule73> | <rule75>
AftAlt9=naval [ <AltList10> ]
25 rule79=3001 [ <AftAlt11> ]
rule80=30 [ <rule81> ]
rule81=01 [ <AftAlt11> ]
rule82=30 [ <rule83> ]
rule83=0 [ <rule84> ]
rule84=1 [ <AftAlt11> ]
30 rule85=3 [ <rule86> ]
rule86=0 [ <rule87> ]
rule87=01 [ <AftAlt11> ]
rule88=3 [ <rule89> ]
rule89=0 [ <rule90> ]
rule90=0 [ <rule91> ]
35 rule91=1 [ <AftAlt11> ]
AltList10=<rule79> | <rule80> | <rule82> | <rule85> | <rule88>
rule93={ :\colon } [ <AftAlt13> ]
rule94={ :\numeric-colon } [ <AftAlt13> ]
AftAlt11=<rule93> | <AftAlt13> | <rule94>
40 AftAlt13=the [ <rule97> ]
rule97=final [ <AltList14> ]
rule98=odyssey [ <AftAlt15> ]
rule99=Odyssey [ <AftAlt15> ]
AltList14=<rule98> | <rule99>
rule101={ .\comma } [ <AftAlt17> ]
45 rule102={ .\numeric-comma } [ <AftAlt17> ]
AftAlt15=<rule101> | <AftAlt17> | <rule102>
AftAlt17=for [ <rule105> ]
rule105=our [ <rule106> ]
rule106=training [ <rule107> ]
rule107=program [ <AltList19> ]
50 rule108={ .\dot } [ <AftAlt19> ]
rule109={ .\period } [ <AftAlt19> ]
rule110={ .\point } [ <AftAlt19> ]
AltList18=<rule108> | <AftAlt19> | <rule109> | <rule110>
AftAlt19=

```

Appendix B

[0089]

5

10

15

20

25

30

35

40

45

50

55

The phoneme label in the first column is pronounced as the capitalized letter(s) in the second field. The phonemes that have a star in the third field are in the subset of 19 used for the rejection grammar.

5

b	Baby		
d	Did		
f	Fife	*	
g	Gag	*	
10	h	Ha	*
j	Jay		
k	Kick	*	
l	Lily	*	
m	Mummy		
15	n	None	*
p	Pop		
r	Rear	*	
s	Sleazy	*	
20	t	Tot	*
v	Vivid	*	
w	We	*	
y	You	*	
z	Zoo		
25	ch	CHip	*
dh	THis		
ng	sING		
sh	SHip	*	
th	THin		
30	zh	meaSure	
L	middle		
M	bottom		
N	kitten		
AA	cat	*	
35	EY	Kate	
AH	calm		
AE	care		
EH	bSt		
eh	Executive		
40	EE	bSt	*
ee	candy		
IH	bIt		
ih	ability		
AY	bIt		
45	OY	boY	
AW	top		
OH	tone	*	
oh	annote		
OW	town		
50	AHH	talk	
UH	luck	*	
uh	Announce		
OOH	too	*	
OC	took		
55	UR	her	
ur	butteR		

Claims

1. A computer-implemented method for enrolling a user in a speech recognition system, comprising:

obtaining data representing a user's speech (1705), the speech including multiple user utterances and corresponding to an enrolment text;

analysing acoustic content of data corresponding to a user utterance (1710);

determining, based on the analysis, whether the user utterance matches a portion of the enrolment text (1730); and

if the user utterance matches a portion of the enrolment text (1730), using the acoustic content of the user utterance to update acoustic models corresponding to the portion of the enrolment text (1735);

wherein determining whether the user utterance matches a portion of the enrolment text comprises permitting a determination that the user utterance matches when the user has skipped words of the enrolment text (1715).

2. The method of Claim 1, wherein obtaining data comprises obtaining data recorded using a recording device physically separate from the computer implementing the method.

3. The method of Claim 2, wherein:

the recording device comprises a digital device, and

obtaining the data comprises receiving a file from the digital recording device.

4. The method of Claim 3, wherein obtaining data comprises receiving signals generated by playing back the user's speech using the recording device.

5. The method of Claim 2, wherein the recording device comprises an analog recording device.

6. The method of any preceding claim, further comprising designating an active portion of the enrolment text (1715), wherein analysing acoustic content of data corresponding to a user utterance comprises analysing the data relative to the active portion of the enrolment text.

7. The method of Claim 6, further comprising identifying a position of a previously analysed utterance in the enrolment text, wherein designating an active portion of the enrolment text comprises designating an active portion based on the identified position.

8. The method of Claim 7, wherein designating the active portion comprises designating a portion that includes text preceding and following the position.

9. The method of Claim 8, wherein designating the active portion comprises designating a portion that includes a paragraph including the position, a paragraph preceding the position, and a paragraph following the position.

10. The method of any one of Claims 6 to 9, wherein analysing the data relative to the active portion of the enrolment text comprises attempting to match the data to models for words included in the active portion of the enrolment text.

11. The method of any one of Claims 6 to 9, wherein analysing the data relative to the active portion of the enrolment text comprises using an enrolment grammar corresponding to the active portion of the enrolment text.

12. The method of any one of Claims 1 to 5, wherein determining whether the user utterance matches a portion of the enrolment text comprises using an enrolment grammar corresponding to the enrolment text.

13. The method of Claim 12, wherein determining whether the user utterance matches a portion of the enrolment text further comprises using a rejection grammar.

14. The method of Claim 13, wherein the rejection grammar comprises a phoneme grammar.

15. The method of Claim 14, wherein the rejection grammar models an utterance using a set of phonemes that is smaller than a set of phonemes used by the enrolment grammar.

16. The method of any preceding claim, further comprising ignoring the user utterance if the user utterance does not match a portion of the enrolment text (1715).

17. A computer-implemented method for enrolling a user in a speech recognition system, comprising:

recording a user's speech using a portable recording device (1705), the speech corresponding to an enrolment text;

transferring the recorded speech to a computer;

using the computer to:

analyse acoustic content of the recorded speech (1710);

identify, based on the analysis, portions of the speech that match portions of the enrolment text (1730);

update acoustic models corresponding to matched portions of the enrolment text using acoustic content of matching portions of the speech (1735).

18. The method of Claim 17, wherein the recorded speech skips or repeats portions of the enrolment text.

19. Computer software, residing on a computer-readable storage medium (145), comprising instructions for causing a computer (125) to:

obtain data representing a user's speech (1705), the speech including multiple user utterances and corresponding to an enrolment text;

analyse acoustic content of data corresponding to a user utterance (1710);

determine, based on the analysis, whether the user utterance matches a portion of the enrolment text (1730); and

use the acoustic content of the user utterance to update acoustic models corresponding to a portion of the enrolment text that matches the user utterance (1735);

wherein the instructions configure the computer to permit a determination that the user utterance matches when the user has skipped words of the enrolment text (1715).

20. A speech recognition system (100) for enrolling a user, comprising:

an input device (105, 165, 170) for receiving speech signals; and

a processor (130) configured to:

obtain data representing a user's speech (1705), the speech (1705) including multiple user utterances and corresponding to an enrolment text;

analyse acoustic content of data corresponding to a user utterance (1710);

determine, based on the analysis, whether the user utterance matches a portion of the enrolment text (1730);

use the acoustic content of the user utterance to update acoustic models corresponding to a portion of

the enrolment text that matches the user utterance (1735); and

permit a determination that the user utterance matches when the user has skipped words of the enrolment text (1715).

Patentansprüche

1. Computer-implementiertes Verfahren zum Registrieren einer Benutzers in einem Spracherkennungssystem, das aufweist:

Erhalten von Daten, die die Sprache eines Benutzers repräsentieren (1705), wobei die Sprache mehrere Benutzeräußerungen enthält und einem Registrierungstext entspricht; Analysieren des akustischen Inhalts von Daten, die einer Benutzeräußerung entsprechen (1710);

Feststellen beruhend auf der Analyse, ob die Benutzeräußerung zu einem Abschnitt des Registrierungstextes paßt (1730); und

wenn die Benutzeräußerung zu einem Abschnitt des Registrierungstextes paßt (1730), Verwenden des akustischen Inhalts der Benutzeräußerung, um akustische Modelle zu aktualisieren, die dem Abschnitt des Registrierungstextes entsprechen (1735);

wobei das Feststellen, ob die Benutzeräußerung zu einem Abschnitt des Registrierungstextes paßt, das Zulassen einer Feststellung aufweist, daß die Benutzeräußerung paßt, wenn der Benutzer Worte des Registrierungstextes ausgelassen hat (1715).

2. Verfahren nach Anspruch 1, wobei das Erhalten von Daten das Erhalten von Daten aufweist, die unter Verwendung einer Aufzeichnungsvorrichtung aufgezeichnet werden, die physikalisch von dem Computer getrennt ist, der das Verfahren implementiert.

3. Verfahren nach Anspruch 2, wobei:

die Aufzeichnungsvorrichtung eine digitale Vorrichtung aufweist, und
das Erhalten der Daten das Empfangen einer Datei von der digitalen Aufzeichnungsvorrichtung aufweist.

4. Verfahren nach Anspruch 3, wobei das Erhalten von Daten das Empfangen von Signalen aufweist, die durch Wiedergeben der Sprache eines Benutzers unter Verwendung der Aufzeichnungsvorrichtung erzeugt werden.

5. Verfahren nach Anspruch 2, wobei die Aufzeichnungsvorrichtung eine analoge Aufzeichnungsvorrichtung aufweist.

6. Verfahren nach einem der vorhergehenden Ansprüche, das ferner die Kennzeichnung eines aktiven Abschnitts des Registrierungstextes aufweist (1715), wobei das Analysieren des akustischen Inhalts von Daten, die einer Benutzeräußerung entsprechen, das Analysieren der Daten relativ zum aktiven Abschnitt des Registrierungstextes aufweist.

7. Verfahren nach Anspruch 6, das ferner das Identifizieren einer Position einer vorhergehenden analysierten Äußerung im Registrierungstext aufweist, wobei die Kennzeichnung eines aktiven Abschnitts des Registrierungstextes die Kennzeichnung eines aktiven Abschnitts beruhend auf der identifizierten Position aufweist.

8. Verfahren nach Anspruch 7, wobei die Kennzeichnung des aktiven Abschnitts die Kennzeichnung eines Abschnitts aufweist, der Text enthält, der der Position vorangeht und ihr folgt.

9. Verfahren nach Anspruch 8, wobei die Kennzeichnung des aktiven Abschnitts die Kennzeichnung eines Abschnitts aufweist, der einen Paragraphen, der die Position enthält, einen Paragraphen, der der Position vorangeht, und einen Paragraphen enthält, der der Position folgt.

10. Verfahren nach einem der Ansprüche 6 bis 9, wobei das Analysieren der Daten relativ zum aktiven Abschnitt des Registrierungstextes das Versuchen aufweist, die Daten an Modelle für Worte anzupassen, die im aktiven Abschnitt des Registrierungstextes enthalten sind.

11. Verfahren nach einem der Ansprüche 6 bis 9, wobei das Analysieren der Daten relativ zum aktiven Abschnitt des Registrierungstextes die Verwendung einer Registrierungsgrammatik aufweist, die dem aktiven Abschnitt des Registrierungstextes entspricht.

12. Verfahren nach einem der Ansprüche 1 bis 5, wobei die Feststellung, ob die Benutzeräußerung zu einem Abschnitt des Registrierungstextes paßt, die Verwendung einer Registrierungsgrammatik aufweist, die dem Registrierungstext entspricht.

13. Verfahren nach Anspruch 12, wobei die Feststellung, ob die Benutzeräußerung zu einem Abschnitt des Registrierungstextes paßt, ferner die Verwendung einer Ablehnungsgrammatik aufweist.

14. Verfahren nach Anspruch 13, wobei die Ablehnungsgrammatik eine Phonemgrammatik aufweist.

15. Verfahren nach Anspruch 14, wobei die Ablehnungsgrammatik eine Äußerung unter Verwendung eines Satzes von Phonemen modelliert, der kleiner als ein Satz von Phonemen ist, der durch die Registrierungsgrammatik verwendet wird.

16. Verfahren nach einem der vorhergehenden Ansprüche, das ferner das Ignorieren der Benutzeräußerung aufweist, wenn die Benutzeräußerung zu keinem Abschnitt des Registrierungstextes paßt (1715).

17. Computer-implementiertes Verfahren zum Registrieren eines Benutzers in einem Spracherkennungssystem, das aufweist:

Aufzeichnen der Sprache eines Benutzers unter Verwendung einer tragbaren Aufzeichnungsvorrichtung (1705), wobei die Sprache einem Registrierungstext entspricht;
Übertragen der aufgezeichneten Sprache zu einem Computer;

wobei der Computer verwendet wird, um:

einen akustischen Inhalt der aufgezeichneten Sprache zu analysieren (1710);
beruhend auf der Analyse Abschnitte der Sprache zu identifizieren, die zu Abschnitten des Registrierungstextes passen (1730);
Aktualisieren akustischer Modelle, die den angepaßten Abschnitten des Registrierungstextes entsprechen, unter Verwendung eines akustischen Inhalts passender Abschnitte der Sprache (1735).

18. Verfahren nach Anspruch 17, wobei die aufgezeichnete Sprache Abschnitte des Registrierungstextes auslöst oder wiederholt.

19. Computer-Software, die auf einem computer-lesbaren Speichermedium (145) liegt, die Befehle aufweist, um einen Computer (125) zu veranlassen:

Daten zu erhalten, die Sprache eines Benutzers repräsentieren (1705), wobei die Sprache mehrere Benutzeräußerungen enthält und einem Registrierungstext entspricht; einen akustischen Inhalt von Daten zu analysieren, die einer Benutzeräußerung entsprechen (1710);
beruhend auf der Analyse zu bestimmen, ob die Benutzeräußerung zu einem Abschnitt des Registrierungstextes paßt (1730); und
den akustischen Inhalt der Benutzeräußerung zu verwenden, um akustische Modelle zu aktualisieren, die einem Abschnitt des Registrierungstextes entsprechen, der zu der Benutzeräußerung (1735);

wobei die Befehle den Computer dazu einrichten, eine Feststellung zuzulassen, daß die Benutzeräußerung paßt, wenn der Benutzer Worte des Registrierungstextes ausgelassen hat (1715).

20. Spracherkennungssystem (100) zum Registrieren eines Benutzers, das aufweist:

eine Eingabevorrichtung (105, 165, 170) zum Empfangen von Sprachsignalen; und
einen Prozessor (130), der dazu eingerichtet ist:

Daten zu erhalten, die die Sprache eines Benutzers repräsentieren (1705), wobei die Sprache (1705)

mehrere Benutzeräußerungen enthält und einem Registrierungstext entspricht;
einen akustischen Inhalt von Daten zu analysieren, die einer Benutzeräußerung entsprechen (1710);
beruhend auf der Analyse zu bestimmen, ob die Benutzeräußerung zu einem Abschnitt des Registrier-
ungstextes paßt (1730);
den akustischen Inhalt der Benutzeräußerung zu verwenden, um akustische Modelle zu aktualisieren, die
einem Abschnitt des Registrierungstextes entsprechen, der zu der Benutzeräußerung paßt (1735); und
eine Feststellung zuzulassen, daß die Benutzeräußerung paßt, wenn der Benutzer Worte des Registrie-
rungstextes ausgelassen hat (1715).

Revendications

1. Procédé mis en oeuvre par ordinateur pour effectuer l'enregistrement d'apprentissage d'un utilisateur dans un système de reconnaissance de la parole, le procédé comprenant les opérations suivantes :

obtenir des données représentant la parole (1705) d'un utilisateur, la parole comportant plusieurs prononcia-
tions de l'utilisation et correspondant à un texte d'enregistrement d'apprentissage ;
analyser le contenu acoustique de données correspondant à une prononciation (1710) de l'utilisateur ;
déterminer, sur la base de l'analyse, si la prononciation de l'utilisateur est appariée à une partie du texte
d'enregistrement d'apprentissage (1730) ; et,
si la prononciation de l'utilisateur est appariée à une partie du texte d'enregistrement d'apprentissage (1730),
utiliser le contenu acoustique de la prononciation de l'utilisateur pour mettre à jour des modèles acoustiques
correspondant à la partie du texte d'enregistrement d'apprentissage (1735) ;

où l'opération de détermination du fait que la prononciation de l'utilisateur est appariée ou non à une partie
du texte d'enregistrement d'apprentissage comprend l'opération consistant à permettre la détermination du fait
que la prononciation de l'utilisateur est appariée lorsque l'utilisateur a sauté des mots du texte d'enregistrement
d'apprentissage (1715).

2. Procédé selon la revendication 1, où l'opération d'obtention de données comprend l'obtention de données enre-
gistrées au moyen d'un dispositif d'enregistrement matériellement séparé, l'ordinateur mettant en oeuvre le pro-
cédé.

3. Procédé selon la revendication 2, où :

le dispositif d'enregistrement comprend un dispositif numérique, et
l'obtention des données comprend l'opération qui consiste à recevoir un fichier de la part du dispositif d'enre-
gistrement numérique.

4. Procédé selon la revendication 3, où l'obtention de données comprend l'opération consistant à recevoir les signaux
produits par reproduction de la parole de l'utilisateur au moyen du dispositif d'enregistrement.

5. Procédé selon la revendication 2, où le dispositif d'enregistrement comprend un dispositif d'enregistrement ana-
logique.

6. Procédé selon l'une quelconque des revendications précédentes, qui comprend en outre l'opération consistant à
désigner une partie active du texte d'enregistrement d'apprentissage (1715), où l'analyse du contenu acoustique
de données correspondant à une prononciation de l'utilisateur comprend l'analyse des données relativement à la
partie active du texte d'enregistrement d'apprentissage.

7. Procédé selon la revendication 6, comprenant en outre l'opération qui consiste à identifier la position d'une pro-
nonciation précédemment analysée du texte d'enregistrement d'apprentissage, où l'opération de désignation d'une
partie active du texte d'enregistrement d'apprentissage comprend la désignation d'une partie active basée sur la
position identifiée.

8. Procédé selon la revendication 7, où l'opération de désignation de la partie active comprend la désignation d'une
partie comportant le texte qui précède et suit la position.

9. Procédé selon la revendication 8, où l'opération de désignation de la partie active comprend la désignation d'une partie qui comporte un paragraphe incluant la position, un paragraphe précédant la position, et un paragraphe suivant la position.

10. Procédé selon l'une quelconque des revendications 6 à 9, où l'opération d'analyse des données relativement à la partie active du texte d'enregistrement d'apprentissage comprend l'opération qui consiste à tenter d'apparier les données à des modèles de mots inclus dans la partie active du texte d'enregistrement d'apprentissage.

11. Procédé selon l'une quelconque des revendications 6 à 9, où l'opération d'analyse des données relativement à la partie active du texte d'enregistrement d'apprentissage comprend l'utilisation d'une grammaire d'enregistrement d'apprentissage correspondant à la partie active du texte d'enregistrement d'apprentissage.

12. Procédé selon l'une quelconque des revendications 1 à 5, où l'opération de détermination du fait que la prononciation de l'utilisateur est appariée à une partie du texte d'enregistrement d'apprentissage comprend l'utilisation d'une grammaire d'enregistrement d'apprentissage correspondant au texte d'enregistrement d'apprentissage.

13. Procédé selon la revendication 12, où l'opération de détermination du fait que la prononciation de l'utilisateur est appariée à une partie du texte d'enregistrement d'apprentissage comprend en outre l'utilisation d'une grammaire de rejet.

14. Procédé selon la revendication 13, où la grammaire de rejet comprend une grammaire de phonèmes.

15. Procédé selon la revendication 14, où la grammaire de rejet modélise une prononciation en utilisant un ensemble de phonèmes qui est plus petit que l'ensemble de phonèmes utilisé par la grammaire d'enregistrement d'apprentissage.

16. Procédé selon l'une quelconque des revendications précédentes, comprenant en outre l'opération qui consiste à ignorer la prononciation de l'utilisateur si la prononciation de l'utilisateur n'est pas appariée à une partie du texte d'enregistrement d'apprentissage (1715).

17. Procédé mis en oeuvre par ordinateur pour effectuer l'enregistrement d'apprentissage d'un utilisateur dans un système de reconnaissance de la parole, comprenant les opérations suivantes :

enregistrer la parole de l'utilisateur au moyen d'un dispositif d'enregistrement portatif (1705), la parole correspondant à un texte d'enregistrement d'apprentissage ;
transférer la parole enregistrée à un ordinateur ;
utiliser l'ordinateur pour :

analyser le contenu acoustique de la parole enregistrée (1710) ;
identifier, sur la base de l'analyse, des parties de la parole qui sont appariées à des parties du texte d'enregistrement d'apprentissage (1730) ;
mettre à jour des modèles acoustiques correspondant à des parties appariées du texte d'enregistrement d'apprentissage en utilisant le contenu acoustique de parties appariées de la parole (1735).

18. Procédé selon la revendication 17, où la parole enregistrée saute ou répète des parties du texte d'enregistrement.

19. Logiciel d'ordinateur, résidant sur un support (145) de stockage pouvant être lu par ordinateur, qui comprend des instructions servant à faire qu'un ordinateur (125) :

obtienne des données représentant la parole (1705) de l'utilisateur, la parole comportant plusieurs prononciations de l'utilisateur et correspondant à un texte d'enregistrement d'apprentissage ;
analyser le contenu acoustique de données correspondant à une prononciation (1710) de l'utilisateur ;
déterminer, sur la base de l'analyse, si la prononciation de l'utilisateur est ou non appariée à une partie du texte d'enregistrement d'apprentissage (1730) ; et
utiliser le contenu acoustique de la prononciation de l'utilisateur pour mettre à jour des modèles acoustiques correspondant à une partie texte d'enregistrement d'apprentissage qui est appariée à la prononciation (1735) de l'utilisateur ;

où les instructions configurent l'ordinateur de manière à permettre la détermination du fait que la prononciation de l'utilisateur est appariée lorsque l'utilisateur a sauté des mots du texte d'enregistrement d'apprentissage (1715).

5 **20.** Système (100) de reconnaissance de la parole destiné à effectuer l'enregistrement d'apprentissage d'un utilisateur, comprenant :

un dispositif d'entrée (105, 165, 170) servant à recevoir des signaux de parole ; et
un processeur (130) configuré pour :

10 obtenir des données représentant la parole (1705) de l'utilisateur, la parole (1705) comportant plusieurs prononciations de l'utilisateur et correspondant à un texte d'enregistrement d'apprentissage;
analyser la teneur acoustique de données correspondant à une prononciation de l'utilisateur (1710) ;
15 déterminer, sur la base de l'analyse, si la prononciation de l'utilisateur est ou non appariée à une partie du texte d'enregistrement d'apprentissage (1730) ;
utiliser le contenu acoustique de la prononciation de l'utilisateur pour mettre à jour des modèles acoustiques correspondant à une partie du texte d'enregistrement d'apprentissage qui est appariée à la prononciation (1735) de l'utilisateur ; et
20 permettre la détermination du fait que la prononciation de l'utilisateur est appariée lorsque l'utilisateur a sauté des mots du texte d'enregistrement d'apprentissage (1715).

25

30

35

40

45

50

55




FIG. 1




FIG. 2




FIG. 3




FIG. 3




FIG. 5





FIG. 6






FIG. 8