

Spoken Language Identification System and Methods for Training and Operating Same

Cross Reference to Related Applications

5 This application claims the benefit of priority to US Provisional Application No. 60/611,022 filed September 17, 2004, the contents of which are herein incorporated by reference in its entirety for all purposes.

Background

10 The present invention relates to language identification systems and methods for training and operating said systems.

 With the growth of globalization, international business, and security considerations, multilingual speech applications are in strong demand, in particular automatic language identification (LID). Possible applications of automatic language
15 identification include automatic call routing, audio mining, and voice automated attendant systems.

 Acoustic-phonotactic based LIDs represent one type of language identification system employed in the art, an illustration of which is shown in Fig. 1. The system typically includes four stages operable to process a speech segment and to classify it
20 into one of a number of possible candidate languages. Initially the system is trained, whereby the system is programmed to recognize particular features of each of the candidate languages. Subsequent to training, language identification operations are performed, whereby a speech sample of unknown language is processed and compared to the previously-programmed features to determine presence or absence of
25 said features, the candidate language possessing the greatest number of correlations with the sample being deemed the language of the sample.

 The conventional system suffers from several disadvantages, one being that a language specific development effort is needed to add a new candidate language. This requirement gives rise to high costs in the acoustic and language modeling and
30 speech data transcription efforts needed. Accordingly, the conventional system is not very scalable with respect to adding new languages.

What is therefore needed is an improved spoken language identification system which provides better scalability with the addition of new candidate languages.

5 Summary

The present invention provides a system and method for identifying spoken language using only a subset of the candidate languages' sound alphabet. Because only a subset of the candidates' sound alphabet is used, speech training and identification operations are faster and more memory efficient.

10 In one representative embodiment, a method for training a spoken language identification system to identify an unknown language as one of a plurality of known candidate languages is presented, the method including the process creating a sound inventory comprising a plurality of sound tokens, the collective plurality of sound tokens provided from a subset of the known candidate languages. The method further
15 includes providing a plurality of training samples, each training sample composed within one of the known candidate languages. Further included is the process of generating one or more training vectors from each training sample, wherein each training vector is defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages. The method further includes
20 associating each training vector with the candidate language of the corresponding training sample.

These and other features of the invention will be better understood when viewed in light of the following drawings and detailed description of exemplary embodiments.

25

Brief Description of the Drawings

Fig. 1 illustrates a spoken language identification system known in the art.

Fig. 2 illustrates a spoken language identification system in accordance with the present invention

30 Fig. 3 illustrates a method for training the spoken language identification system in accordance with the present invention.

Fig. 4 illustrates a method for identifying the language of an unknown language sample in accordance with the present invention.

Fig. 5A illustrate classification error rates of the spoken language identification system provide in accordance with the present invention.

5 Fig. 5B illustrate the classification error rates of the spoken language identification system of the present invention as compared to a prior art acoustic-phonotactic system.

Fig. 6 illustrates exemplary classification error rates as a function of training data for the spoken language identification system of the present invention.

10 For clarity previously identified features retain their reference numerals in subsequent drawings.

Detailed Description of Exemplary Embodiments

15 Language Identification System

Fig. 2 illustrates a spoken language identification system in accordance with the present invention. The system 200 includes a sound recognizer 210 and a sound classifier 220. Optionally, the system 200 may include a database 230 for storing training samples 232_{1-n} , and a database 240 for storing training vectors 242_{1-n} .

20 Alternatively, databases 230 and 240 may be located remotely from the system 200.

The sound recognizer 200 includes an input coupled to receive training samples 230_{1-n} , the sound recognizer 200 also operable to receive (either via the same input or another input) an unknown language sample 250. The reader will appreciate that the training and/or unknown language samples may be provided in several
25 different forms, e.g., as live or recorded sound, the speech itself being in any format, e.g., analog or digital forms.

Each training sample 230_i includes one or more predefined speech utterances composed in one of the candidate languages which are to be identified by the system. The one or more training utterances defining each training sample 230_i are selected to
30 provide a majority of the phonemes or other phonotactical markers employed in that particular candidate language. Each training sample may employ any number of training utterances of any duration, an exemplary embodiment being between 100-

3000 separate training utterances employed, each speech utterance being between 5 and 15 seconds long. Further exemplary, the unknown speech sample 250 is of a similar duration as the speech utterances, for example, between 5-15 seconds in duration.

5 The sound recognizer 210 is operable to generate a training vector 240_i from a received training sample 230_i during system operation in the language training mode, and to generate an unknown language vector 260 from the unknown language sample 250 during operation in the identification mode. The generated training and unknown language vectors 240_i and 260 (both described below) are supplied to database 240
10 and the sound classifier 220, respectively, as shown. In an alternative embodiment of the invention, the sound recognizer 210 is configured to generate training vectors for each one or combination of the speech utterances included within each training sample 230_i . In such an instance, the resulting training vectors corresponding to the different speech utterances in the same training sample are collectively processed to
15 form a statistical model used in the language identification process, as will be further described below.

 In both the training and identification modes of operation, the sound recognizer 210 generates the training and unknown language vectors 240_i and 260 as a function of phonemes 212 which are either supplied to the sound recognizer 210 or
20 stored locally within the sound recognizer 210. In a particular embodiment, the phonemes 212 are obtained from a subset of the n candidate languages which the system 200 is operable to identify. A small set of phonemes obtained from a subset of the n candidate languages can be used to identify a larger number of languages due to the observation that many languages employ the same phonemes, albeit differently in
25 a phonotactical sense. In the present invention, the sound recognizer 210 generates vectors 240_i and 260 which model the phonotactical characteristics of the supplied training or unknown language samples, using as its alphabet, the smaller phoneme set, that set including at least some phonemes which are reusable for two or more of the candidate languages. It is to be noted that while aforementioned generation of
30 the training and unknown language vectors is described in terms of phonemes, it will be understood that any phonotactical marker may be used. In addition, an acoustically steady-state sound token may be used. Further, a phonetic subword such

as syllable may be used as a sound token. The term “sound token” is used to describe each of these parameters.

The system 200 further includes a sound classifier 220 coupled to receive training vectors 240_i during training mode operations, as well as an unknown language vector 260 during speech identification mode operations. The sound classifier is operable to identify the language of the unknown speech sample 250 as the language associated with the training vector exhibiting the closest correlation to the unknown language vector 260, a process of which is further detailed below.

Method of System Training

Fig. 3 illustrates a method for training the spoken language identification system 200 in accordance with the present invention. Initially at 302, a sound inventory (e.g., 212 in Fig. 2) is created with sound tokens that are taken from a subset of the candidate languages the system 200 is operable to identify. In an exemplary embodiment illustrated below, the sound inventory is composed of sound tokens from three languages, while the speech identification system is operable to identify six different candidate languages. Other arrangements are of course possible, for example, a system operable to identify multiple candidate languages using sound tokens from only a single one of them. As noted above, the sound inventory 212 may be formed as a part of the system 200, or located remotely therefrom.

At 304, training samples 230_i composed within each of the candidate languages are provided. The training samples 230_i may be provided locally within the system or supplied from an external source.

Next at 306, one or more training vectors 240_i are generated from each of the training samples 230_i , the training vectors generated as a function of the aforementioned sound tokens. For a system of V sound tokens, in one embodiment of this process, the training vectors 240_i are composed of counts of sound tokens in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$, wherein n is the candidate language index, L is the length of the sound token sequence. Therefore, each training vector 240_i has $D=V$ dimension, each element represents the count of each sound token. Let v be v^{th} sound token in the said sound inventory and c_v^n be the number of occurrence of said

v^{th} sound token in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$ divided by the total number of occurrence of all the sound tokens.

$$c_v^n = \frac{c(v)}{\sum_v c(v)}$$

In another embodiment of the invention, the training vectors 240_i are composed of counts of sound token pairs in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$. As a sound inventory of V sound tokens lead to $N = V \times V$ sound token pairs, each training vector 240_i has $N = V \times V$ dimension, each element represents the count of each sound token pair $\{v_j v_k\}$.

$$c_{v_j v_k}^n = \frac{c(v_j v_k)}{\sum_{v_j v_k} c(v_j v_k)}$$

Another embodiments of this invention include larger phonotactic units such as sound token triplets, quartets. Those skilled in the art will appreciate that other embodiments are also possible in the present invention.

In a particular embodiment, process 306 is performed as a part of the sound recognizer processes in one embodiment of the invention described above. Further as noted above, each of the training samples 230_i may include several speech utterances, each (or a combination of two or more) of which is used as the basis to generate a training vector 240_i, the collection of which may be used to form a statistical model for determining language identity of the unknown language sample. At 308, each training vector 240_i is associated with the candidate language of the corresponding training sample 230. The collection of training vectors 240_n associated with language n are used to represent the nth language. Further, a statistical model λ_n is constructed for each language n based on the collection of training vectors 240_n. In this way, a collection of training vectors are represented by a statistical model. In an exemplary embodiment of this process, the arithmetic mean of the collection of training vectors is used to form a mean vector, said mean vector forming the basis of the statistical model λ_n .

Method of Language Identification

Fig. 4 illustrates a method for identifying the language of an unknown language sample in accordance with the present invention. The process largely builds upon the aforementioned training process in which the system 200 is trained to recognize n candidate languages using an alphabet of phonotactic markers taken from a subset of the candidate languages.

In the exemplary embodiment shown, the process begins at 402 in which an unknown language sample is received and an unknown language vector is generated therefrom, the unknown language vector defined as a function of the aforementioned sound tokens taken from a subset of the candidate languages. Exemplary embodiments of the unknown language vector include the vector of sound tokens or vector of sound token pairs, sound token triplets and quartets as described above.

At 404, the unknown language vector is correlated to one or more of the training vectors 240, wherein the language associated with the training vector having the closest correlation to the unknown language vector is identified as the language of the unknown language sample. In a particular embodiment of process 404, the unknown language vector X is extracted, and used to correlate with each of the training vector statistical models λ_n to determine correlation between the unknown language vector and language n :

$$\hat{n} = \arg \max_n P(X / \lambda_n)$$

with the statistical model exhibiting the closest correlation (exhibiting the highest probability) being deemed the optimal statistical model, and its corresponding language being adjudged the language of the unknown language sample.

In a further embodiment of the invention, frequency in terms of count of the sound tokens and their combinations in different orders such as pairs, triplets and quartets can be recorded and arranged in the language vectors. Implementation of such vectors enables a global assessment and correlation between the unknown language vector and the training vectors, or statistical model resulting from the training vectors, rather than the shorter duration analysis used in phone n -gram language models. The present invention not only incorporates the statistics of local sound token co-occurrences by using sound token pairs, triplets and quartets as the

elements of language vector as in the prior art, but also incorporates the statistics of long distance, separated sound token in which sounds can be separated, for example, by 10-20 seconds and with hundreds of intervening sounds.

5 Exemplary System Embodiment & Performance

 An exemplary implementation of the language identification system 200 was constructed to identifying six candidate languages Chinese Mandarin, Chinese Cantonese, Chinese Shanghai dialect, Japanese, Korean and English using phonemes from only three of these languages Chinese Mandarin, English and Korean. The
10 system 200 implemented a pre-processing prior to the sound recognizer, including extracting, for each speech frame, a 39-dimensional feature vector consisting of 12 Mel Frequency Cepstral Coefficients and normalized energy, plus the first and second order derivatives. Sentence-based cepstral mean subtraction is applied for acoustic normalization in both in the training and identification modes. The 39 dimensional
15 vector is subsequently supplied to the sound recognizer 210.

 A training corpus of training samples 230 was used to train the sound inventory 210 required for later system use in the language identification mode. Typically about 10 hours of speech is required for training. The specific embodiment tested used telephony speech input digitized with a sampling rate of 8 kHz. It used
20 124 phonemes consisting of 43, 44 and 37 phonemes respectively from combining Chinese Mandarin, English and Korean into a training corpus. Consistent with the terminology in the application, each of these phonemes is called a “sound token”. As noted above, the inventory of basic "sound token" can be defined as a set of phonemes, acoustically steady sound tokens, and the like.

25 Each sound token was modeled by a left-to-right three-emitting-state CDHMM, each state having 32 Gaussian mixture components. In addition to the 124 sound tokens emitted from these states, multiple noise models were built to absorb noise events of different kinds. A set of flat models based on a broad “sound” definition is used as the reference ground to obtain the acoustic confidence score. An
30 adaptive voice detector was used to detect the start and end point of the speech. The sound classifier 220 is implemented as a support vector machine (SVM), although in other implementation, latent semantic analysis (LSA), an artificial neural network, or

any other high dimensional vector classifier may be used. During language identification operations, the sound recognizer 210 converts the input speech of unknown language into sound sequences. The sound classifier first extracts the unknown language vector from the sequences, and then measures the unknown language vector against all the training vectors of interests, the language whose training vector gives rise to the highest similarity scores to the unknown language vector, was identified as the result language. The length of the time window used in the sound classifier is varied from 5 to 10 to 15 seconds. The number of training exemplars is varied from 100 to 3000.

Fig. 5A illustrate classification error rates of the spoken language identification system as reported in "Spoken Language Identification Using Bag of Sounds," authors Ma Bin and Li Haizhou (March 21, 2005, International Conference on Chinese Computing 2005), herein incorporated by reference. For each language, separate databases for training and testing the sound classifier 220 was constructed, and the training databases was not used in testing.

For each test set at 5, 10 and 15 seconds, 2,000 training vectors were used to build the SVM classifier 220. The testing corpus included 500 unknown language vectors of 5, 10 and 15 seconds. Fig. 5A shows that longer input utterances lead to much better classification accuracy, which is as expected. The average accuracy for six languages in the 10-second test is 98.1% in accuracy or 1.9% in error rate.

Fig. 5B illustrate the classification error rates of the spoken language identification system (identified as BOS) as compared to a prior art acoustic-phonotactic system (identified as AP) as reported in the aforementioned article to Ma Bin. The test included the three baseline languages, Chinese Mandarin, English and Korean. Three mono-lingual phoneme recognizers were constructed as well as their respective phoneme bi-gram language models. A phone loop grammar was used in acoustic decoding, and the phoneme bi-gram was used to derive final phoneme sequences. An unknown language sample is provided through three parallel acoustic-phonotactic decoders, the winner of combined acoustic and language model scores takes all. The greatly improved accuracy provided by the present invention over the prior art acoustic-phonotactic system is readily apparent.

Fig. 6 illustrates exemplary classification error rates as a function of the number of training samples for the spoken language identification system of the present invention. The number of training samples was varied from 100 to 3,000 for each training sample of 5, 10 and 15 seconds. The data indicates that fewer training samples of longer duration are needed. For example, only 100 15 second duration training sessions are needed to provide an acceptable error rate of approximately 2 percent, whereas approximately 2000 ten second training samples are needed to achieve the same accuracy.

As readily appreciated by those skilled in the art, the described processes may be implemented in hardware, software, firmware or a combination of these implementations as appropriate. In addition, some or all of the described processes may be implemented as computer readable instruction code resident on a computer readable medium (removable disk, volatile or non-volatile memory, embedded processors, etc.), the instruction code operable to program a computer or other such programmable device to carry out the intended functions.

The foregoing description has been presented for purposes of illustration and description. It is not intended to be exhaustive or to limit the invention to the precise form disclosed, and obviously many modifications and variations are possible in light of the disclosed teaching. The described embodiments were chosen in order to best explain the principles of the invention and its practical application to thereby enable others skilled in the art to best utilize the invention in various embodiments and with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the claims appended hereto.

CLAIMS

What is claimed is:

- 5 1. A method of training a spoken language identification system to identify an unknown language as one of a plurality of known candidate languages, the method comprising:

creating a sound inventory comprising a plurality of sound tokens, the collective plurality of sound tokens provided from a subset of the known candidate
10 languages;

providing a plurality of training samples, each training sample composed within one of the known candidate languages;

generating one or more training vectors from each training sample, wherein each training vector is defined as a function of said plurality of sound tokens provided
15 from said subset of the known candidate languages, wherein each element of the training vector is a function of the respective sound token according to a sound token sequence; and

associating each training vector with the candidate language of the corresponding training sample.

20

2. The method of claim 1, wherein the sound inventory comprises a plurality of V sound tokens, the training vector comprising counts of sound tokens in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$, wherein n is the candidate language index, L is the length of the sound token sequence, wherein said sound token count comprises c_v^n ,
25 where v is defined as the v^{th} sound token in the said sound inventory and c_v^n is defined as the number of occurrences of said v^{th} sound token in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$ divided by the total number of occurrences of all the sound tokens:

$$c_v^n = \frac{c(v)}{\sum_v c(v)}$$

3. The method of claim 1, wherein the sound inventory comprises a plurality of V sound tokens, the training vector comprising counts of sound token pairs in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$. wherein n is the candidate language index, L is the length of the sound token sequence, wherein said count of sound token pairs
5 comprises $c_{v_j v_k}^n$, where v_j, v_k is defined as the number of occurrences of the sound token pairs in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$ divided by the total number of occurrences of all the sound token pairs $\{v_j v_k\}$:

$$c_{v_j v_k}^n = \frac{c(v_j v_k)}{\sum_{v_j v_k} c(v_j v_k)}$$

10

4. The method of claim 1, wherein the sound inventory comprises a plurality of sound tokens, and wherein the training vector comprises counts of sound token triplets.
- 15 5. The method of claim 1, wherein the sound inventory comprises a plurality of sound tokens, and wherein the training vector comprises counts of sound token quartets.
6. The method of any one of claims 1-5, wherein generating one or more training
20 vectors comprises generating a plurality of training vectors corresponding to one candidate language, the method further comprising constructing a statistical model from the plurality of the training vectors associated with the candidate language.
7. The method of claim 6, wherein the arithmetic mean of said plurality of
25 training vectors is used to form a mean training vector, said mean training vector forming the basis of the statistical model λ_n used to determine the correlation between said mean training vector and an unknown training vector associated with an unknown language sample.

8. In a spoken language identification system operable to identify the language of a unknown language sample as one or a plurality of known candidate languages, the language identification system including (i) means for creating a sound inventory comprising a plurality of sound tokens, the collective plurality of sound tokens
5 provided from a subset of the known candidate languages, (ii) means for providing a training sample composed within each of the known candidate languages, (iii) means for generating one or more training vectors from each training sample, wherein each training vector is defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages, wherein each element of the
10 training vector is a function of the respective sound token according to a sound token sequence, and (iv) means for associating each training vector with the candidate language of the corresponding training sample, a method for identifying the language of the unknown speech sample, comprising
generating an unknown language vector from the unknown language sample,
15 the unknown language vector defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages; and
correlating the unknown language vector to a plurality of the training vectors,
wherein the language associated with the training vector having the closest correlation to the unknown language vector is identified as the language of the of the
20 unknown language sample.

9. The method of the claim 8, wherein one or more training vectors corresponding to the candidate language n are represented by a statistical model λ_n , and wherein correlating comprises correlating the unknown language vector X to one
25 or more of the training vector statistical models λ to determine correlation between the unknown language vector and language n :

$$\hat{n} = \arg \max_n P(X / \lambda_n)$$

whereby, the statistical model exhibiting the highest probability being deemed the optimal statistical model, and its corresponding language being adjudged the language of the unknown language sample.

5 10. A spoken language identification system operable to identify an unknown language as one of a plurality of known candidate languages, the system comprising:

a sound recognizer coupled to receive an unknown language sample, the sound recognizer operable to generate an unknown language vector from the unknown language sample, said unknown language vector being defined as a function
10 of a plurality of sound symbols, wherein the plurality of sound symbol are provided from a subset of the candidate languages; and

a sound classifier coupled to receive the unknown language vector, the sound classifier operable to correlate the unknown language vector with one or more predefined training vectors, wherein each element of the training vector is a function
15 of the respective sound token according to a sound token sequence, each having a candidate language associated therewith, the sound classifier operable to identity the language of the unknown language sample as the language of the training vector having the closest correlation with the unknown language vector.

20 11. The spoken language identification system of claim 10, wherein the sound recognizer is further operable to receive a plurality training samples, each of the training samples composed in one of candidate languages, the sound recognizer further comprising:

a sound inventory comprising said plurality of sound tokens provided from a
25 subset of the candidate languages; and

means for generating a candidate language vector from the training sample, the candidate language vector being associated with one of the candidate languages, and the candidate language vector defined as a function of said plurality of sound symbols provided from a subset of the candidate languages.

30

12. The system of one of claims 10 or 11, wherein the sound inventory comprises a plurality of V sound tokens, each training vector comprising counts of sound tokens

in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$, wherein n is the candidate language index, L is the length of the sound token sequence, wherein said sound token count comprises c_v^n , where v is defined as the v^{th} sound token in the said sound inventory and c_v^n is defined as the number of occurrences of said v^{th} sound token in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$ divided by the total number of occurrences of all the sound tokens:

$$c_v^n = \frac{c(v)}{\sum_v c(v)}$$

13. The system of one of claims 10 or 11, wherein the sound inventory comprises a plurality of V sound tokens, each training vector comprising counts of sound token pairs in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$. wherein n is the candidate language index, L is the length of the sound token sequence, wherein said count of sound token pairs comprises $c_{v_j v_k}^n$, where v_j, v_k is defined as the number of occurrences of the sound token pairs in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$ divided by the total number of occurrences of all the sound token pairs $\{v_j v_k\}$:

$$c_{v_j v_k}^n = \frac{c(v_j v_k)}{\sum_{v_j v_k} c(v_j v_k)}$$

14. The system of one of claims 10 or 11, wherein the sound inventory comprises a plurality of sound tokens, and wherein the training vector comprises counts of sound token triplets.

15. The system of one of claims 10 or 11, wherein the sound inventory comprises a plurality of sound tokens, and wherein the training vector comprises counts of sound token quartets.

16. A computer program product, resident on a computer readable medium, operable to control a programmable system to train a spoken language identification system to identify an unknown language as one of a plurality of known candidate languages, the computer program product comprising:

5 instruction code to create a sound inventory comprising a plurality of sound tokens, the collective plurality of sound tokens provided from a subset of the known candidate languages;

instruction code to provide a plurality of training samples, each training sample composed within one of the known candidate languages;

10 instruction code to generate one or more training vectors from each training sample, wherein each training vector is defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages, wherein each element of the training vector is a function of the respective sound token according to a sound token sequence; and

15 instruction code to associate each training vector with the candidate language of the corresponding training sample.

17. The computer program product of claim 16, wherein the sound inventory comprises a plurality of V sound tokens, each training vector comprising counts of
20 sound tokens in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$, wherein n is the candidate language index, L is the length of the sound token sequence, wherein said sound token count comprises c_v^n , where v is defined as the v^{th} sound token in the said sound inventory and c_v^n is defined as the number of occurrences of said v^{th} sound token in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$ divided by the total number of
25 occurrences of all the sound tokens:

$$c_v^n = \frac{c(v)}{\sum_v c(v)}$$

18. The computer program product of claim 16, wherein the sound inventory
30 comprises a plurality of V sound tokens, each training vector comprising counts of

sound token pairs in the sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$. wherein n is the candidate language index, L is the length of the sound token sequence, wherein said count of sound token pairs comprises $c_{v_j v_k}^n$, where v_j, v_k is defined as the number of occurrences of the sound token pairs in said sound token sequence $m_1^n, \dots, m_l^n, \dots, m_L^n$
5 divided by the total number of occurrences of all the sound token pairs $\{v_j v_k\}$:

$$c_{v_j v_k}^n = \frac{c(v_j v_k)}{\sum_{v_j v_k} c(v_j v_k)}$$

19. The computer program product of claim 16, wherein the sound inventory
10 comprises a plurality of sound tokens, and wherein the training vector comprises counts of sound token triplets.

20. The computer program product of claim 16, wherein the sound inventory
15 comprises a plurality of sound tokens, and wherein the training vector comprises counts of sound token quartets.

21. The computer program product of any one of claims 16-20, wherein the instruction code to generate one or more training vectors comprises instruction code to generate a plurality of training vectors corresponding to one candidate language,
20 the computer program product further comprising instruction code to construct a statistical model from the plurality of the training vectors associated with the candidate language.

22. The computer program product of claim 21, wherein the arithmetic mean of
25 said plurality of training vectors is used to form a mean training vector, said mean training vector forming the basis of the statistical model λ_n used to determine the correlation between said mean training vector and an unknown training vector associated with an unknown language sample.

23. In a spoken language identification system operable to identify the language of a unknown language sample as one or a plurality of known candidate languages, the language identification system including (i) means for creating a sound inventory comprising a plurality of sound tokens, the collective plurality of sound tokens
5 provided from a subset of the known candidate languages, (ii) means for providing a training sample composed within each of the known candidate languages, (iii) means for generating one or more training vectors from each training sample, wherein each training vector is defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages, wherein each element of the
10 training vector is a function of the respective sound token according to a sound token sequence, and (iv) means for associating each training vector with the candidate language of the corresponding training sample, a computer program product, resident on a computer-readable medium, operable to control a programmable system to identify the language of the unknown speech sample supplied thereto, the computer
15 program product comprising:

instruction code to generate an unknown language vector from the unknown language sample unknown language, the unknown language vector defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages; and

20 instruction code to correlate the unknown language vector to a plurality of the training vectors,

wherein the language associated with the training vector having the closest correlation to the unknown language vector is identified as the language of the of the unknown language sample.

25

24. The computer program product of the claim 23, wherein one or more training vectors corresponding to the candidate language n are represented by a statistical model λ_n , and wherein the instruction code to correlate comprises instruction code to correlate the unknown language vector X to one or more of the training vector
30 statistical models λ to determine correlation between the unknown language vector and language n :

$$\hat{n} = \arg \max_n P(X / \lambda_n)$$

whereby, the statistical model exhibiting the highest probability being deemed the optimal statistical model, and its corresponding language being adjudged the language of the unknown language sample.

25. A method of operating a spoken language identification system to identify an unknown language as one of a plurality of known candidate languages, the method comprising:

10 creating a sound inventory comprising a plurality of sound tokens, the collective plurality of sound tokens provided from a subset of the known candidate languages;

providing a plurality of training samples, each training sample composed within one of the known candidate languages;

15 generating one or more training vectors from each training sample, wherein each training vector is defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages, wherein each element of the training vector is a function of the respective sound token according to a sound token sequence;

20 associating each training vector with the candidate language of the corresponding training sample;

generating an unknown language vector from the unknown language sample, the unknown language vector defined as a function of said plurality of sound tokens provided from said subset of the known candidate languages; and

25 correlating the unknown language vector to a plurality of the training vectors, wherein the language associated with the training vector having the closest correlation to the unknown language vector is identified as the language of the of the unknown language sample.

30 26. The method of the claim 25, wherein one or more training vectors corresponding to the candidate language n are represented by a statistical model λ_n , and wherein correlating comprises correlating the unknown language vector X to one

or more of the training vector statistical models λ to determine correlation between the unknown language vector and language n :

$$\hat{n} = \arg \max_n P(X / \lambda_n)$$

5

whereby, the statistical model exhibiting the highest probability being deemed the optimal statistical model, and its corresponding language being adjudged the language of the unknown language sample.

10