



(51) International Patent Classification:

G10L 17/10 (2013.01) G10L 15/22 (2006.01)
G10L 17/00 (2013.01)

(21) International Application Number:

PCT/GB2016/051954

(22) International Filing Date:

29 June 2016 (29.06.2016)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

62/346,036 06 June 2016 (06.06.2016) US

(71) Applicant: **CIRRUS LOGIC INTERNATIONAL SEMICONDUCTOR LIMITED** [GB/GB]; 7B Nightingale Way, Quartermile, Edinburgh EH3 9EG (GB).

(72) Inventors: **AVILÉS-CASCO VAQUERO, Carlos**; C/ Constanca 27 1ºE, 28002 Madrid (ES). **GARCÍA GOMAR, Marta**; Calle Bola 12 Piso Bajo A, 28013 Madrid

(ES). **MARTÍNEZ GONZÁLEZ, David**; Avda. Cesareo Alierta 47, 3B, esc.1, 50008 Zaragoza (ES).

(74) Agent: **O'CONNELL, David**; Haseltine Lake LLP, Redcliff Quay, 120 Redcliff Street, Bristol BS1 6HU (GB).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,

(54) Title: VOICE USER INTERFACE




Figure 9

(57) **Abstract:** A received signal represents a user's speech. A first speaker recognition process is performed on a first portion of the received signal, to obtain a first output result. A second speaker recognition process is performed on a second portion of the received signal that is different from the first portion of the received signal, to obtain a second output result. The second speaker recognition process is different from the first speaker recognition process. The first and second output results are combined to obtain a combined output result indicating a likelihood that the user is a registered user.



TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

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

VOICE USER INTERFACE

Technical Field

- 5 The embodiments described herein relate to a method and system for use in a voice user interface, for example for allowing a user to control the operation of a device using speech.

10 Background of the invention

Voice user interfaces are provided to allow a user to interact with a system using their voice. One advantage of this, for example in devices such as smartphones, tablet
15 computers and the like, is that it allows the user to operate the device in a hands-free manner.

In one typical system, the user wakes the voice user interface from a low-power standby mode by speaking a trigger phrase. Speech recognition techniques are used to detect that the trigger phrase has been spoken and, separately, a speaker
20 recognition process is used to confirm that the trigger phrase was spoken by a registered user of the device.

The voice user interface may then provide a prompt to the user, to confirm that the system is active, and the user may then speak a command, which can be recognised
25 by the voice user interface using speech recognition techniques.

The voice user interface may then act on that spoken command. For example, if the spoken command asks for publicly available information, the spoken command may be recognised, and used to generate a query to an internet search engine in order to be
30 able to supply that information to the user.

However, in other cases, for example if the spoken command relates to personal information, the level of authentication provided by the speaker recognition process may be considered insufficient for the voice user interface to act on that command. In
35 such cases, the user may be asked to provide an additional form of authentication, for

example by entering a PIN number or password through a keypad of the device, or by providing additional biometric authentication, such as a fingerprint scan.

5 This means that the user is no longer able to operate the device in a hands-free manner.

Summary of the invention

10 According to the embodiments described herein, there is provided a method and a system which reduce or avoid one or more of the disadvantages mentioned above.

According to a first aspect of the invention, there is provided a method of processing a received signal representing a user's speech, the method comprising: performing a first
15 speaker recognition process on a first portion of the received signal, to obtain a first output result; performing a second speaker recognition process on a second portion of the received signal that is different from the first portion of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from the first speaker recognition process; and combining the first and second output
20 results to obtain a combined output result indicating a likelihood that the user is a registered user.

According to another aspect of the invention, there is provided a device for processing a received signal representing a user's speech, for performing speaker recognition,
25 wherein the device is configured to: perform a first speaker recognition process on a first portion of the received signal, to obtain a first output result; perform a second speaker recognition process on a second portion of the received signal that is different from the first portion of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from the first speaker recognition
30 process; and combine the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.

According to a further aspect of the invention, there is provided an integrated circuit device for processing a received signal representing a user's speech, for performing
35 speaker recognition, wherein the integrated circuit device is configured to: perform a first speaker recognition process on a first portion of the received signal, to obtain a

first output result; perform a second speaker recognition process on a second portion of the received signal that is different from the first portion of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from the first speaker recognition process; and combine the first and second
5 output results to obtain a combined output result indicating a likelihood that the user is a registered user.

According to another aspect of the invention, there is provided a method of processing a received signal representing a user's speech, the method comprising: performing a
10 first speaker recognition process on a first portion of the received signal, to obtain a first output result; performing a second speaker recognition process on a second portion of the received signal, to obtain a second output result; combining the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user; determining if the indicated likelihood that the user is a
15 registered user is higher than a threshold value; performing a speech recognition process on the second portion of the received signal to obtain a command contained in the second portion of the received signal; and acting on the obtained command if the indicated likelihood is higher than the threshold value.

20 In another aspect, the present invention provides a device for processing a received signal representing a user's speech, for performing speaker recognition, wherein the device is configured to: perform a first speaker recognition process on a first portion of the received signal, to obtain a first output result; perform a second speaker recognition process on a second portion of the received signal that is different from the first portion
25 of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from the first speaker recognition process; and combine the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.

30 The invention also provides an integrated circuit device for processing a received signal representing a user's speech, for performing speaker recognition, wherein the integrated circuit device is configured to: perform a first speaker recognition process on a first portion of the received signal, to obtain a first output result; perform a second speaker recognition process on a second portion of the received signal that is different
35 from the first portion of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from the first speaker recognition

process; and combine the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.

5 The invention also provides a non-transitory computer readable storage medium having computer-executable instructions stored thereon that, when executed by processor circuitry, cause the processor circuitry to perform any of the methods set out above.

10 Brief description of the drawings

For a better understanding of the invention, and to show more clearly how it may be carried into effect, reference will now be made, by way of example only, to the accompanying drawings in which:

15

Figure 1 is a schematic view of an electronic device;

Figure 2 is a further schematic diagram of an electronic device;

20 Figure 3 is a flow chart, illustrating a method;

Figure 4 is a plot illustrating a step in the method;

Figure 5 is a schematic diagram illustrating a step in the method;

25

Figure 6 is a plot illustrating a step in the method;

Figure 7 is a plot illustrating a step in the method;

30 Figure 8 is a plot illustrating a further step in the method;

Figure 9 is a schematic diagram illustrating a first method; and

Figure 10 is a schematic diagram illustrating a second method.

35

Detailed description

For clarity, it will be noted here that this description refers to speaker recognition and to speech recognition, which are intended to have different meanings. Speaker
5 recognition refers to a technique that provides information about the identity of a person speaking. For example, speaker recognition may determine the identity of a speaker, from amongst a group of previously registered individuals, or may provide information indicating whether a speaker is or is not a particular individual, for the purposes of identification or authentication. Speech recognition refers to a technique
10 for determining the content and/or the meaning of what is spoken, rather than recognising the person speaking.

Figure 1 shows a device in accordance with one aspect of the invention. The device may be any suitable type of device, such as a tablet computer, a games console, a
15 remote control device, a domestic appliance including a domestic temperature or lighting control system, a toy, a machine such as a robot, or the like, but in this illustrative example the device is a smartphone 10. The smartphone 10 may, by suitable software, be used as the control interface for controlling any other further device or system.

20

The smartphone 10 includes a screen 12 for displaying information to a user, a sound inlet 14, for allowing sound to be detected by a microphone, and a jack socket 16, or other port or receptacle, for allowing an accessory to be connected to the device.

25 **Figure 2** is a schematic diagram showing the smartphone 10. In this example, the smartphone 10 includes a microphone 20, which may for example be located close to the sound inlet 14 shown in Figure 1. Electronic signals generated by the microphone 20 are passed to a signal processing block 22, which performs initial signal processing of the signals, for example converting analog signals received from the microphone 20
30 into digital signals.

The smartphone 10 also includes an accessory interface 24, which may for example be located close to the jack socket 16 shown in Figure 1. The jack socket 16 and the interface 24 may be suitable for allowing a headset accessory to be connected thereto,
35 and signals received from a microphone on such an accessory are also passed to the signal processing block 22, which performs initial signal processing of the signals.

The signal processing block 22 is connected to a processor 26, which performs methods as described herein on the basis of data and program instructions stored in a memory 28.

5

The processor 28 is connected to an interface 30, which is itself connected to an antenna 32, allowing signals to be transmitted and received over an external network to remote devices.

10 In other examples, the device performing the processes described herein may receive the required input signals in a suitable form, without needing to perform any prior signal detection or signal processing and thus not requiring the device to comprise signal processing block 22.

15 In some examples, some of the processing described below may be performed on an external device communicated with via an external network, for example a remote computing server or a server in a home network. In other examples, all of the processing described below may be performed in a single device, without requiring the device to comprise any interface to any external device or network.

20

Figure 3 is a flow chart, illustrating a method of operation of a voice user interface according to one embodiment.

As described in more detail below, the process shown in Figure 3 is performed after a user has registered with the system, for example by providing one or more sets of voice samples that can be used to form one or more model of the user's speech.

Typically, the registration or enrolment process requires the user to provide speech inputs, and then uses these speech inputs to form a model of the user's speech, starting from a particular background model defined in a prior development phase.

30 Thus, the background model and the speech inputs are the inputs to the enrolment process that is used to form the model of the user's speech. Subsequently, during verification, as described in more detail below, further speech inputs are compared with the model of the user's speech, and with a background model, in order to provide an output. The output may for example be a numerical value indicating a likelihood that
35 the speech inputs received during the verification phase were provided by the same user that provided the speech inputs during enrolment. The numerical value indicative

of the likelihood may be for example a log likelihood ratio (LLR) or may be some more indirect indication, for example a metric of distance of extracted features of the speech sample from some one- or multi-dimensional threshold.

5 The voice user interface may spend the majority of its time in a standby state, in order to save power. A voice activity detection block may be provided, for determining when sounds that are detected by a microphone represent speech. In some embodiments, signals that are received from a microphone are stored. Then, when the voice activity block determines that the sounds that are detected by the microphone represent
10 speech, those stored signals are analysed as described below.

In step 50, a signal is received from one or more microphone, and the signal represents a first voice segment, that is, the signal represents the first part of a user's speech.

15 In step 52, the signal representing the first voice segment is passed to a low power trigger phrase detector, which may for example be implemented in a processor (such as the processor 26 in the embodiment shown in Figure 2) within the device, or may be implemented in a separate device, for example a computing server in a home or
20 remote network.

The trigger phrase detector uses speech recognition techniques for determining whether the first voice segment contains a specific predetermined phrase, referred to herein as a trigger phrase or pass phrase. This is a phrase that the user is required to
25 speak, in order to fully activate the voice user interface.

If the trigger phrase detector fails to recognise the trigger phrase at step 52, the process returns to step 50, in which the system continues monitoring the received signal for voice segments.

30 If it is determined at step 52 that the first voice segment contains the trigger phrase, the process passes to step 54, in which a speaker recognition process is performed on the first voice segment.

35 The speaker recognition process that is performed on the first voice segment operates with a first user model. Because the user is expected to be speaking a trigger phrase

that is known in advance, the first user model may use text-constrained, or text-dependent, speaker recognition. That is, during an enrolment process performed by the user when first activating the speaker recognition feature of the device, the user was required to speak the trigger phrase multiple times. The enrolment process uses
5 as its starting point a particular background model, which may for example be based upon the speech inputs from a large sample of other users, and the user's speech inputs are then used in forming the first user model. More specifically, features of the user's speech, such as Mel frequency Cepstrum Coefficient (MFCC) features may be used in forming the first user model.

10

The output of step 54 may for example be a numerical score (in the form of a likelihood ratio, or a distance, or in any other form). This numerical score relating to the first voice segment, or trigger phrase, is denoted here by s_T .

15 In step 56, a further signal is received from the one or more microphone, and this signal represents a second voice segment, that is, the signal represents the second part of a user's speech. This second voice segment may follow immediately on from the first voice segment. Alternatively, the system may provide a visual or audible prompt to the user on detecting the trigger phrase at step 52, and the second voice
20 segment may then follow that prompt.

In step 58, a speaker recognition process is performed on the second voice segment.

The speaker recognition process that is performed on the second voice segment differs
25 from the speaker recognition process that was performed on the first voice segment in step 54.

Each speaker recognition process uses a particular background model and a model of the user's speech as its inputs, and compares the relevant voice segment with these
30 models, using a specified verification method to arrive at the output. Features of the user's speech, such as Mel frequency Cepstrum Coefficient (MFCC) features, are obtained from the relevant voice segment, and these features are compared with the features of the background model and the relevant user model. Thus, each speaker recognition process can be considered to comprise the background model, the user
35 model, and the verification method or engine that are used. The speaker recognition

processes performed on the first and second voice segments may differ in one or more of these components of the process.

For example, the speaker recognition processes may use the same model of the user's speech (for example, by using just one enrolment session and technique), but may use different verification methods for the first and second voice segments. As one example, the background model and user model may be Hidden Markov Models (HMMs), with the background model being a Universal Background Model (UBM) adapted to human speech, and the user model being adapted to that specific user's speech, but not adapted to any phrase in particular. Then the Text Constrained method used on the trigger phrase may use a forced alignment method to force the HMMs to go through the path indicated by the trigger phrase (a fixed sequence of states in the HMM as defined in the UBM), while the Text Independent system used on the second or command phrase may look for the sequence giving the best alignment without forcing anything. Another example is the case where the UBM and user model are Gaussian Mixture Models (GMMs). The same GMM models can be used in a simple GMM-UBM system that will be fast and low power, for example for the trigger phrase, and then used in a more complex GMM system for example incorporating Factor Analysis, that considers the same speaker models and UBM (and that needs more computational cost and power, but obtains much more accuracy), for the command.

Alternatively, the methods may use different models of the user's speech (also referred to as "user models"), which implies that the methods use different enrolment sessions or processes.

One way to obtain different user models is to have different engines operating on the same set of input audio data during the enrolment process. That is, input audio data is used to form a user model during an enrolment process, as described above. In practice, this is achieved by applying the input audio data to a set of mathematical equations that have certain associated parameters as constants in the equations. Different engines, which may for example be tuned for different purposes, can be obtained by using a different set of equations, including by setting different values for these parameters or by employing a different class of algorithm. For example the input could be a GMM-UBM that goes to an enrolment process to get a user adapted GMM,

or to an enrolment process that turns the GMM-UBM into a supervector and builds a Support Vector Machine (SVM) as user model.

Another way to obtain different user models for the speaker recognition processes performed on the first and second voice segments is to use different input data but a common engine during the enrolment process. A third way to obtain different user models for the speaker recognition processes performed on the first and second voice segments is to use different input data during the enrolment process and to use different engines operating on the input data.

When it is desired to obtain different user models for the speaker recognition processes performed on the first and second voice segments by using different input data, one possibility is to use two different audio inputs during the enrolment process to train the respective speaker models. For example, the audio inputs used to form the user model that is used in the speaker recognition process performed on the first voice segment may be audio inputs in which the user speaks the predetermined trigger phrase, so that this becomes a text-constrained speaker recognition process, while the audio inputs used to form the user model that is used in the speaker recognition process performed on the second voice segment may be audio inputs in which the user speaks (or reads aloud) any phrases of their choosing (or a predefined text or set of phrases that is designed to have high phonetic variability), so that this becomes a text-independent speaker recognition process.

A second possibility is to use different background models in order to form the different user models that are used in the speaker recognition processes performed on the first and second voice segments. One specific possibility is to use two different Gaussian Mixture Models (GMMs) as the background models, and hence in the derived user model, by using different numbers of Gaussian components in the models. This may be useful because reducing the number of Gaussian components makes the process faster, while increasing the number of Gaussian components makes the process more accurate.

A third possibility for obtaining the different user models for the speaker recognition processes performed on the first and second voice segments is to use different audio inputs during the enrolment process and to use these to modify respective different background models.

The output of step 58 may for example be a numerical score (in the form of a likelihood ratio, or a distance, or in any other form). This numerical score relating to the second voice segment, which may for example be a command, is denoted here by s_c .

5

In step 60, the results of the speaker recognition processes performed on the first and second voice segments are combined, in order to obtain a combined output result indicating a likelihood that the user is the enrolled user. The combined output result indicative of the likelihood may be for example a log likelihood ratio (LLR) or may be
10 some more indirect indication, for example a metric of distance of extracted features of the speech sample from some one- or multi-dimensional threshold or nominal point or volume in a multi-dimensional speech parameter space.

The combined output result may be obtained from the separate results of the speaker
15 recognition processes performed on the first and second voice segments by any suitable method. For example, the combined output result may be a weighted sum s_f of the results of the speaker recognition processes performed on the first and second voice segments. That is, in general terms:

20

$$s_f = \alpha s_T + \beta s_c + \gamma$$

The weighting factors α , β , and γ may be constant and determined in advance.

Alternatively, the step of combining the results of the speaker recognition processes
25 performed on the first and second voice segments, to obtain a combined output result, may use quality measures to determine how the results should be combined, in order to improve the reliability of the decision. That is, separate quality measures are obtained for the first and second voice segments, and these quality measures are then used as further inputs to the process by which the results are combined.

30

These quality measures may for example be based on properties of the first and second voice segments, for example the trigger phrase and the command. Certain triggers will be more suitable for use in voice biometrics than others because they are longer in duration, or because they contain more phonetic variability and thus they
35 provide more information to differentiate speakers. Certain commands will be more suitable for use in voice biometrics than others for the same reasons. Other aspects,

such as the presence of non-stationary noise in either the first and second voice segments may make one voice segment more reliable than the other.

In one embodiment there is defined a set of quality measures, namely a set of quality
 5 measures Q_T for the trigger and a set of quality measures Q_C for the command, and the values of weighting factors α , β , and γ are set based on the quality measures. Then a weighted sum s_F will be obtained as a function of these quality measures:

$$s_F = \alpha(Q_T, Q_C)s_T + \beta(Q_T, Q_C)s_C + \gamma(Q_T, Q_C)$$

10

The functions that map the quality measures Q_T , Q_C to the weighting factors α , β , and γ are part of the system design and are thus obtained and defined during a development phase, before the system is deployed for user enrolment or verification. The values
 15 returned by these functions in use after the development phase will vary from sample to sample as the quality measures Q_T , Q_C vary from sample to sample.

20

The functions may be obtained during the development phase on the basis of exercising the system with a large number of speech samples arranged to have a range of different values of the quality measures.

The form of the functions may be defined before the development phase, and coefficients or optimised to provide the best fit. In some embodiments, the functions may not be algebraic functions but may comprise the form of a look-up table containing
 25 optimised coefficients optimised over ranges of value of the quality measures or fixed values applied to optimised ranges of quality measures. More generally a function may be the result of some more complex algorithm characterised by some coefficients and delivering a value dependent on the quality measures.

In some embodiments the combined score may be a non-linear combination of the
 30 scores s_F and s_C , which may for example be represented in the form

$$s_F = \alpha(Q_T, Q_C, s_T)s_T + \beta(Q_T, Q_C, s_C)s_C + \gamma(Q_T, Q_C)$$

where the each weighting factor α or β may depend continuously or non-continuously
 35 on the respective score.

More generally, the combined score may be any function of the scores, s_T and s_C , that are obtained from the speaker recognition processes performed on the first and second voice segments, and of the quality measures, Q_T and Q_C , that apply to those voice segments. That is:

5
$$s_F = f(s_T, s_C, Q_T, Q_C)$$

where f may be any function.

The values of the scores, s_T and s_C , and of the quality measures, Q_T and Q_C , may be applied to a neural network, which then produces a value for the combined score s_F

10

When determining the weights to be given to the results of the first and second speaker recognition processes, different quality measures can be considered.

One suitable quality measure is the Signal to Noise Ratio (SNR), which may for
15 example be measured in the input trigger and in the input command separately. In the case of non-stationary noise, where the SNR varies rapidly, a higher weight can be given to the result obtained from the input speech segment that has the higher SNR.

Another suitable quality measure is the net-speech measure. As discussed in
20 connection with the illustrated embodiment, the weight that is given to the score obtained from the command can be increased according to the amount of speech in the command. That is, the total length of the fragments in the command that actually contain speech, excluding non-speech segments, is measured, for example in time units such as seconds, and this is used to form the weight to be applied to the
25 command, relative to the weight applied to the trigger.

A further alternative quality measure is a saturation measure. This determines the amount of the audio input that is saturated, this is, where part of the input audio is beyond the dynamic range of the sensor and the pre-processing stages that process
30 the audio before it reaches the speaker recognition system. This usually happens when the input level is too loud for the system. Saturation is a non-linear effect that generates unexpected harmonics in the signal, and reduces significantly the speaker recognition accuracy since all signals affected by saturation are similar from the speaker recognition point of view, and very distinct from non-saturated signals. Thus, it is
35 possible to measure the amount of audio input that has been saturated in the trigger and in the command. Then, a higher weight is given to the less saturated portion of

audio. Saturation can be measured in several ways, but one typical way of measuring it is indicating the percentage of the frames (the number of windows of a certain length, for example 10 ms) that contains saturated audio samples. This percentage is usually calculated over those frames that are labelled as speech by the system (that is, those frames that will be processed by the speaker recognition system).

A further alternative type of quality measure is a measure related to how well the input matches the model, that is, the level of fit between the input and the background models. Universal Background Models (UBM) for speaker recognition are typically statistical models that try to model the universe of inputs that are expected from the universe of speakers. It is expected that the UBM is a good model to describe any input that may be fed into the speaker recognition system. If the UBM is not a close model to an input, we expect the decision of the speaker recognition system to be less reliable, since the input is not similar to anything the system has seen before.

One example of this type of measure is a measure of the likelihood of the data for the text-dependent and text-independent UBM. If one segment of the input has the expected likelihood values (for example those observed on a large development dataset) and the other segment has values below the expected likelihood values, more weight is given to the former system, since it is more capable of describing its input data.

Another example of a quality measure related to how well the input matches the model concerns a divergence between the prior and posterior values of hidden variables. Thus, HMM, GMM and Factor analysis systems use hidden or latent variables in order to explain the input data to the system. When treated in a Bayesian way, there is a prior on these latent variables. Given the input data, a posterior can be obtained. If the divergence between the prior and posterior is large, this means that the input data is not similar to what the model would expect and the system output may not be reliable. This information, measured in both the trigger and the command, can be used to adjust the weights that are given to the results obtained from the speaker recognition processes performed on the trigger and on the command.

A further alternative type of quality measure is a measure provided by an external module of the device that works with the voice biometrics solution.

For example, a measure of this type of may be obtained where the device includes a beamformer. A beamformer may provide measures of the Direct to Reflected Ratio, for each segment of the input signal. If one segment of the input signal (for example the first voice segment that might be the trigger of the second voice segment that might be the command) seems less affected by reverberation or indirect reflections, the system may increase the weight of that segment. This variation may for example happen when the relative positions of the user and the device vary. The system may thus give more weight to a portion of speech according to a measure of a direction from which a respective portion of speech was received, in order to weight more heavily a signal portion more directly received from the user. Somewhat similarly, a beamformer may provide measures of the Signal to Interference ratio of the received signals. Then, when there is more than one source of sounds (for example the desired speaker and some directional noise or other interfering speakers), the system can give more weight to the segments that are less affected by the interferers.

A further measure of this type may be obtained where the device includes a proximity sensor. For example, a proximity sensor can determine how far the user was from the microphone when speaking the first and second voice segments, i.e. the range from which the respective portion of the received signal was received. This information can be used to give more weight to the speaker recognition results obtained on voice segments that were uttered from a shorter distance.

During the verification phase, in one embodiment the method may detect the start and the end of the trigger and of the command, and then extract the quality measures and the scores for the trigger and the command, using the corresponding background models and user models obtained during enrolment, and these background models and/or user models may be for the trigger and the command. The quality measures are used to determine the weighting factors, based on the functions learnt during the development stage. Then the weighting factors are used to obtain a combined result from the separate results obtained by performing the first and second speaker recognition processes on the trigger phrase and the command respectively.

Particularly, in one embodiment, the weighting factors α , β and γ may be determined based on the particular trigger phrase considered and the amount of user speech available in the command, n_c . For this purpose, a development phase is needed.

During the development phase, optimal values are obtained for the weighting factors α , β , γ are obtained for different lengths of net-speech n_c (that is, the amount of actual user speech) in the command, considering always the same trigger phrase. Once the trigger is fixed, the optimal weighting factors for a given value of n_c , that is $\alpha(n_c)$, $\beta(n_c)$, $\gamma(n_c)$ may be selected according to certain criteria. For example, Linear Logistic Regression can be used to determine the optimal values of the weighting factors.

Linear Logistic Regression applies Logistic Regression to the linear combination of the scores to obtain the combined score s_F , thus:

10

$$s_F = \alpha(n_c)s_T + \beta(n_c)s_C + \gamma(n_c)$$

The Logistic Regression tries to determine the true value that the decision θ of the system should take ($\theta = 1$ for hypothesis H_1 , that is a match, and $\theta = 0$ for hypothesis H_2 , that is no-match) from the score through the logistic function:

15

$$P(H_1 | s_T, s_C, n_c) = \frac{1}{1 + e^{-(\alpha(n_c)s_T + \beta(n_c)s_C + \gamma(n_c))}}$$

The output of the Logistic Function can be interpreted as the probability that the input segment is a match to the speaker model, given the scores for the trigger and the command, and the net-speech length of the command. Note that, in a Bayesian framework, this is only true for a flat prior, that is $P(H_1) = P(H_2) = 0.5$. However, for different priors, it is still possible to include prior information in the Logistic Function.

20

Figure 4 shows how this sigmoid function $P(H_1 | s_F)$, that is $P(H_1 | s_T, s_C, n_c)$, asymptotically approaches 1 or 0 as the combined score s_F increases or decreases.

25

Thus:

in the limit, as $(s_F \rightarrow -\infty)$, $P(H_1 | s_F) = 0$ and

30

in the limit, as $(s_F \rightarrow \infty)$, $P(H_1 | s_F) = 1$.

Scaling or biasing the inputs to this function (as altering the weighting factors will in general do), will shape the sigmoid accordingly. Specifically, changing the weighting factors $\alpha(n_c)$ and $\beta(n_c)$ will shape the function so that the slope will be more or less steep, and changing the bias weighting factor $\gamma(n_c)$ will move the function along the horizontal axis.

35

In order to determine the optimal values for the weighting factors, or fusion coefficients, the development process needs, as training data, a large number (for example in the thousands) of samples of scores of the text-constrained and text-independent speaker
 5 recognition systems for a particular n_c and trigger phrase, and the decisions associated with these scores. Every training sample i is a set of values $\{s_T, s_C, n_C, \theta\}_i$, where $\theta = 1$ for a match, and $\theta = 0$ for a no-match. This set has been extracted from an actual user interaction with the VUI. Then Linear Logistic Regression determines the optimal set of coefficient values that best predicts the true value of θ with $P(H_1 | s_T, s_C, n_C)$.

10

Specifically, this is achieved by minimizing the cross-entropy between the true underlying distribution of the output decision $P(\theta)$ (observed decisions) and the model $P(H_1 | s_T, s_C, n_C)$. Note that the cross-entropy minimization process does not set any constraint in the fusion weights, and so their values are not bounded.

15

This process is done for several n_C values, so that a set of fusion coefficients is obtained for each n_C . In practice, it is more useful to agglomerate the samples that have similar values for n_C , rather than those that have exactly the same value of n_C . For this purpose, K intervals of n_C values can be defined, and the input samples can be
 20 pre-classified according to the interval that their n_C falls into.

Figure 5 illustrates this development process. As mentioned above, there are a large number of training samples 160, and these are passed to a pre-classifier, where they are classified based on the interval that their value of n_C falls into. Thus, $(K-1)$
 25 thresholds are defined, and for values of n_C that are less than a first threshold n_1 , the samples are passed to a first linear logistic regression block 164.1; for values of n_C that are less than a first threshold n_k , but greater than the previous threshold n_{k-1} , the samples are passed to a k^{th} linear logistic regression block 164.k; and for values of n_C that are greater than the last threshold n_{K-1} , the samples are passed to a K^{th} linear
 30 logistic regression block 164.K.

Then, as shown at 166, a function is carried out, designed to map every value of n_C , including those seen and not seen during development, to obtain the corresponding fusion coefficients 168. Finally, a calibration operation 170 is performed, so that the
 35 threshold 172 is adjusted to the desired operating point. Although this step can alternatively be included in the Linear Logistic Regression, this separate calibration

enables the setting of constraints to the fusion weight values, facilitating their modelling.

A simple non-linear mapping using step functions could be considered, so that for each interval of n_c values considered for pre-classification, a set of fusion coefficients is defined as follows:

$$\begin{aligned} \{\alpha(n_c), \beta(n_c), \gamma(n_c)\} &= \{\alpha_1, \beta_1, \gamma_1\} & n_c < n_1 \\ \{\alpha(n_c), \beta(n_c), \gamma(n_c)\} &= \{\alpha_k, \beta_k, \gamma_k\} & n_{(k-1)} \leq n_c < n_k \\ \{\alpha(n_c), \beta(n_c), \gamma(n_c)\} &= \{\alpha_K, \beta_K, \gamma_K\} & n_{(K-1)} \leq n_c \end{aligned}$$

An example of this piece-wise constant mapping is presented in the table below, where six intervals have been defined, noting that $n_c = 0$ is a special case, in which there is no command.

n_c	$\alpha(n_c)$	$\beta(n_c)$	$\gamma(n_c)$
0	1.00	0.00	0.00
$0 < n_c < 1$	2.89	1.50	8.87
$1 < n_c < 2$	2.83	2.07	10.24
$2 < n_c < 3$	2.76	2.45	11.17
$3 < n_c < 5$	2.73	2.62	11.82
$5 < n_c < 10$	2.71	2.80	12.38
$10 < n_c$	2.75	3.18	13.85

A regression model can be used for the mapping as well. The use of a regression model to estimate the fusion weights from the net-speech is more natural, since net-speech is a continuous variable rather than a discrete one.

The fusion weights obtained from Logistic regression do not have any constraint, and so regression models may not fit very well. The weights may therefore be scaled and biased in such a way as to obtain a better fit.

Figure 6 is a plot of the values in the table above, showing the weights applied to the score obtained from the text-constrained (TC) speaker recognition process (i.e. the weight $\alpha(n_c)$ applied to the score obtained from the command) and to the score

obtained from the text-independent (TI) speaker recognition process (i.e. the weight $\beta(n_c)$ applied to the score obtained from the trigger), for the six ranges of values of the net-speech present in the command (i.e. n_c). For each of these ranges, the scores are shown at the average value of the net-speech for all of the inputs in the respective

5 range for the particular set of samples actually used to generate these weight values. As would be expected, the weight applied to the score obtained from the command increases relative to the weight applied to the score obtained from the trigger, as the amount of net-speech present in the command increases.

10 To facilitate the task of building a regression model in step 166 of Figure 5, it is possible to set a constraint to the weights. The offset term $\gamma(n_c)$ in the fusion weights translates into a movement of the threshold that is used to distinguish between a match condition and a no-match condition. Since a calibration stage (shown at 170 in Figure

15 5) is used to set the desired threshold, this offset term is not needed.

Also, the scale of the fusion weights $\alpha(n_c)$, $\beta(n_c)$ is only important for calibration, as far as the relationship $\alpha(n_c)/\beta(n_c)$ is maintained. Therefore, ignoring the scale of $\alpha(n_c)$ and $\beta(n_c)$ and ignoring the offset term $\gamma(n_c)$, it is possible to set a constraint for the weights, namely:

20
$$\alpha(n_c) + \beta(n_c) = 1$$

Figure 7 is another plot of the values from the table above, showing the weights applied to the score obtained from the text-constrained (TC) speaker recognition process (i.e. $\alpha(n_c)$) and to the score obtained from the text-independent (TI) speaker

25 recognition process (i.e. $\beta(n_c)$), for the six ranges of values of the net-speech present in the command (i.e. n_c), after applying the constraint that $\alpha(n_c) + \beta(n_c) = 1$, while maintaining the ratio $\alpha(n_c)/\beta(n_c)$.

This constraint makes the values of $\alpha(n_c)$ and $\beta(n_c)$ more reasonable, since the

30 tendency to increase of one translates in the tendency to decrease of the other, avoiding ranges of n_c where both increase or decrease, which make difficult modelling these weights.

It is possible to use a regression model for the ratio $\alpha(n_c)/\beta(n_c)$, rather than for each

35 weight separately. In one embodiment, a linear regression model is used to estimate the ratio $\alpha(n_c)/\beta(n_c)$ for any given value of n_c . In fact, a specific embodiment, uses a

model that is linear with respect to n_c^{-1} , since we expect asymptotic behaviour for very large values of n_c . Thus, values of parameters a and b are obtained that provide the best fit to a line of the form:

$$\alpha(n_c)/\beta(n_c) = a.n_c^{-1} + b$$

Figure 8 is a plot of the values from the table above, converted into the form required by this regression equation. That is, for each of the ranges of values of n_c for which values of $\alpha(n_c)$ and $\beta(n_c)$ were obtained, the average value of the ratio $\alpha(n_c)/\beta(n_c)$ is plotted against n_c^{-1} , the reciprocal of the average value of the net-speech for all of the inputs in the respective range. Then, a best-fit straight line is found, which in this case gives values of parameters $a = 0.7289$ and $b = 0.855$, and therefore:

$$\alpha(n_c)/\beta(n_c) = 0.7289.n_c^{-1} + 0.855.$$

In this illustrated case, the values were obtained for a particular set of triggers. In one alternative, separate values can be obtained for each trigger, and the regression can be performed separately for each trigger, leading to different terms in the regression for each trigger.

Then, for any future input values of n_c that are obtained from received voice commands, the linear model can be used to find a value for the ratio $\alpha(n_c)/\beta(n_c)$, and then values for the separate fusion weights $\alpha(n_c)$ and $\beta(n_c)$ can be calculated using the constraint that $\alpha(n_c) + \beta(n_c) = 1$.

As an alternative to this linear regression technique, a non-linear regression techniques, for example as the map to a step function such as that previously presented, or a more complex non-linear function, could be used.

As a further alternative to the process described above, the values of n_c could be included directly in the Linear Logistic Regression process, avoiding the need for performing a pre-classification. In this case, the Logistic Function $\sigma()$ is applied to n_c before inputting it to the Linear Logistic Regression, so that $\sigma(c.n_c + d)$ takes values between 0 and 1, and these values are used to weight the input samples, for the Linear Logistic Regression. The output model of the logistic regression then is capable of providing the fusion coefficients for any input value n_c . This is known as using “side

information” in the Linear Logistic Regression. This also eliminates the need for a final calibration stage, although in practice it may be preferable to perform a final calibration to obtain the threshold value.

- 5 In addition to the determination of the fusion coefficients, during development, the text-constrained speaker recognition system may be adapted to the trigger phrase using some or all of the same samples of the trigger phrase.

10 In addition, or alternatively, the trigger detection system may be adapted to the trigger phrase.

During enrolment, the user may provide samples of the same trigger phrase used in development for creating the user voiceprint for the text-constrained voice biometrics system, and may also provide text-independent samples for creating the user
15 voiceprint for the text-independent voice biometrics system. The samples of the trigger phrase may also be used to adapt and tune the trigger detector.

So far, it has been assumed that the text-constrained and text-independent voice biometrics engines that operate on the first and second voice segments operate
20 independently. However, they may share information to provide more robust responses. For example, in an embodiment in which the engines that operate on the first and second voice segments are identical, except for the background models that they use, the text-independent engine operating on the second voice segment may use partial or complete information from the first voice segment (i.e. the trigger) to obtain a
25 more accurate score

So far, it has been implied that the text-constrained and text-independent voice biometrics engines operate on first and second voice segments which are non-overlapping. However in some embodiments it may be advantageous to operate the
30 later check on a second voice segment which comprises some or all of the first voice segment, i.e. the two voice segments analysed are overlapping rather than non-overlapping. For instance, the first voice segment may be analysed using a relatively fast and simple speaker recognition engine, whereas the second voice segment may be analysed using a more complex but more reliable speaker recognition algorithm,
35 and the result of the latter may be made even more reliable by increasing the effective

length of the sample analysed. Thus in some embodiments the first and second voice segments may be different but overlapping.

In some embodiments a further authentication process may be performed in parallel
5 with the trigger or command speaker authentication, and the corresponding authentication results may be combined in a similar fashion to the above.

This parallel authentication process may be an authentication process based on different characteristics of speech, for example a time-domain based method. In some
10 embodiments the parallel authentication process may be one sensitive to spectral or other characteristics symptomatic of an attempt to spoof the system by playback or direct injection techniques, i.e. be an antispoofing technique.

An exemplary antispoofing technique is based on computing a likelihood ratio (LR). It is
15 extracted from a feature vector, y computed from the testing signal, and two Gaussian models $N(y; \mu_{(\text{non-})\text{spoof}}, \Sigma_{(\text{non-})\text{spoof}})$, one that represents non-spoof data and other that represents spoof data:

$$LR = \frac{N(y; \mu_{\text{spoof}}, \Sigma_{\text{spoof}})}{N(y; \mu_{\text{non-spoof}}, \Sigma_{\text{non-spoof}})} \quad (1)$$

20

where μ_{spoof} and Σ_{spoof} are the mean vector and the diagonal covariance matrix for the spoof model, and $\mu_{\text{non-spoof}}$ and $\Sigma_{\text{non-spoof}}$ are the mean vector and the diagonal covariance matrix for the non-spoof model, which are typically also speaker dependent and they are typically obtained during enrolment.

25

The antispoofing feature vector is composed of different metrics, for example by the spectral ratio, low frequency ratio and feature vector squared Mahalanobis distance. In this example, $N_{AS}=3$. Alternatively other metrics could also be included or one or two of the metrics may be replaced by other metrics or omitted entirely.

30

The spectral ratio may for example be the ratio between the signal energy from 0 to 2 kHz and from 2 kHz to 4 kHz. Thus, given a frame l of the audio access $x(n)$, the spectral ratio may be calculated as:

$$SR(l) = \sum_{f=0}^{\frac{NFFT}{2}-1} 20 \log_{10}(|X(f,l)|) \sqrt{\frac{4}{NFFT}} \cos\left(\frac{(2f+1)\pi}{NFFT}\right) \quad (2)$$

where $X(f,l)$ is the Fast Fourier Transform (FFT) value of the frame l and the f frequency bin of the audio signal, and $NFFT$ is the number of points of FFT (256 samples, for example).

After computing eq. (2) for all the frames, the average value of the spectral ratios (SR_{audio}) may be calculated as the mean of the spectral ratios of the frames whose modulation index is above a given threshold (for example 0.75). Frames with a modulation index above a given threshold usually correspond to speech signals, such that typically by calculating the mean of the spectral ratios of frames with a modulation index above a given threshold results in calculating the mean describing the parts of the audio signal comprising speech. The frames can be created from the audio signal using a window length of 20 msec. with shift of 10 msec, for example.

15

The low frequency ratio may e.g. be computed as the ratio between the signal energy from 100 Hz to 300 Hz and from 300 Hz to 500 Hz. Given a frame l , it may be calculated as:

$$LFR(l) = \sum_{f=100\text{Hz}}^{300\text{Hz}} 20 \log_{10}(|X(f,l)|) - \sum_{f=300\text{Hz}}^{500\text{Hz}} 20 \log_{10}(|X(f,l)|) \quad (3)$$

After computing eq. (3) for all the frames, the average value of the spectral ratios (LFR_{audio}) may be calculated as the mean of the low frequency ratios of the frames whose modulation index is above a given threshold (for example 0.75). The frames can be created using a window length of 20 msec. with shift of 10 msec, for example.

Finally, the feature vector squared Mahalanobis distance may be computed between the average MFCC vectors (e.g. with dimension $N_{\text{AS-MFCC}}$) along the time obtained from the speaker recognition testing signal received in the audio access(es), and the one observed during enrolment. A standard deviation diagonal matrix is typically also required to compute the distance (which is typically computed during or after enrolment).

30

A standard deviation diagonal matrix may e.g. be computed as disclosed in the European patent application EP-2860706A. An alternative, improved way of using the result of an Antispoofing process is described here. In this example, the voice first and second segments, which may be the trigger and the command as discussed previously, are subject to separate antispoofing detection processes (which may be the same or different), to obtain two antispoofing output scores, one for the trigger and one for the command.

These scores may then be integrated with the speaker recognition scores.

One way of using the antispoofing output scores is to use them as a filter.

In this case, where we have two different antispoofing scores (one for the trigger r_T and one for the command r_C), the system may reject the user when one of them is over a very high threshold ($r_T > \epsilon_{TH}$ OR $r_C > \epsilon_{CH}$), suggesting that it is very likely that the input trigger or command is a spoof. The system may also reject the user if both the trigger and the command antispoofing scores are over certain lower threshold ($r_T > \epsilon_{TL}$ AND $r_C > \epsilon_{CL}$, with $\epsilon_{TH} > \epsilon_{TL}$ and $\epsilon_{CH} > \epsilon_{CL}$).

An alternative would be to apply the same fusion method described in this application for speaker recognition to fuse the antispoofing scores and reject the user if the resulting score r_F is over a certain threshold:

$$r_F = \alpha_{AS} \cdot r_T + \beta_{AS} \cdot r_C + \gamma_{AS}.$$

Note that, in this case, the fusion weights will not be the same as those obtained for speaker recognition, but they can be obtained as function of quality measures, in the same way as explained for speaker recognition.

A further option is to use the antispoofing scores as a quality measure in the speaker recognition process described previously. In this case the fusion weights for speaker recognition will be modified according to the antispoofing output for the trigger and the command. The antispoofing scores are thus used in a similar manner to the net-speech in the previously-described embodiment, or as we would use any other quality measure. In general it is a good idea to combine this approach with the use of the

antispoofing output scores as a filter, so that, if the input is clearly a spoof it is directly rejected instead of using the antispoofing scores only to modify the weights.

A third option is to use the antispoofing scores as additional scores for the speaker recognition task, and fuse them with the speaker recognition scores. This approach is advantageous if we expect there to be correlation between the speaker recognition and antispoofing scores. In this case, before fusing the speaker recognition scores for the trigger and the command, we fuse the trigger speaker recognition and antispoofing scores, and we fuse the command speaker recognition and antispoofing scores, obtained as output of one or more speaker recognition systems and one or more antispoofing systems (for simplicity the equations below are presented considering only one speaker recognition system and one antispoofing system for trigger and one speaker recognition system and one antispoofing system for the command), so that the new scores for trigger and the command are:

15

$$p_T = f_T(s_T, r_T)$$

$$p_C = f_C(s_C, r_C)$$

The trigger and command fusion functions may be identical or different, and they may be a linear fusion, which in general is the most reasonable approach when the inputs are assumed to be uncorrelated, or any other fusion method. Particularly, the trigger and command fusion score could be obtained using a classifier that is capable of modelling the correlations that may appear between the multiple speaker recognition and anti-spoofing scores. To do so, for each trigger input or command input, a score vector can be built using each score as an entry of the vector, for example for the trigger and the case of having one speaker recognition and one anti-spoofing score:

25

$$\begin{bmatrix} s_T \\ r_T \end{bmatrix}.$$

For N scores, this vector is a point in an N dimensional space. Using several thousand of samples as development data, a classifier can be trained to properly classify the input vectors of scores into the two classes of interest: user (so the output of the system should be a match) or impostor/spoof. This classifier may be a Gaussian classifier with two classes (this is sometimes known as Gaussian Back End), or may use more complicated models, as GMMs. It also may be a SVM or other type of classifier.

35

In general the classifier will output a new score on how likely is the input vector of scores to belong to the class of the user, or ultimately, the input to be a match. This score may be again in the form of a likelihood ratio, a distance or something else, and
 5 this score will be used directly used as p_T (in the case of the trigger) for the later fusion of trigger and command.

Exactly the same process would be applied for the command, using $\begin{bmatrix} s_c \\ r_c \end{bmatrix}$ as the score
 vector and obtaining p_c . Note that the score vectors considered for the command may
 10 contain the same number of scores as the score vectors for the trigger, or a different number of scores (i.e., antispoofing may be used for the trigger only, or a different number of speaker recognition systems and anti-spoofing systems may be used for the trigger and the command).

15 This third option can also be combined with the use of the antispoofing output scores as a filter, as described above.

The use of the antispoofing scores as a quality measure can also be combined with the use of the antispoofing scores as additional scores for the speaker recognition task,
 20 particularly if the systems that are used to obtain quality measures are not considered for obtaining the scores p_T and p_c , and vice-versa.

Thus, various embodiments have been described above, for arriving at a combined result of the speaker recognition.

25 In step 62, it is determined based on this combined result whether the user is authenticated. For example, it may be determined that the user is authenticated if the combined result obtained as a function of the separate results exceeds a threshold value.

30 **Figure 9** is a schematic diagram, illustrating a first embodiment of the authentication process that is generally as described above.

Thus, a first background model 110 for use in a text-constrained (TC) speaker
 35 recognition process and a second background model 112 for use in a text-independent

(TI) speaker recognition process are taken as starting points. During an initial enrolment, user speech samples are used to form a text-constrained user voiceprint (that is, the set of coefficients that define the user model) 114 and a text-independent user voiceprint 116.

5

A trigger detection block 118 detects that a trigger phrase has been spoken and a text-constrained speaker recognition or voice biometrics process 120 is performed, based on the first voice segment (namely the user speech from the trigger phrase), the first background model 110, and the text-constrained user voiceprint 114, to produce a

10 score s_T . Also, a text-independent speaker recognition or voice biometrics process 122 is performed, based on the second voice segment (namely the user speech received after the trigger phrase, which is expected to be a command), the second background model 112, and the text-independent user voiceprint 116, to produce a score s_C .

15 A first set of quality measures (Q_T) 124 is obtained from the first voice segment, and a second set of quality measures (Q_C) 126 is obtained from the second voice segment.

The functions that map the quality measures Q_T , Q_C to the weighting factors α , β , and γ have been obtained during a development phase and are stored at 128. Thus, the

20 values of the quality measures Q_T , Q_C are used to obtain values of the weighting factors $\alpha(Q_T, Q_C)$, $\beta(Q_T, Q_C)$, and $\gamma(Q_T, Q_C)$ as shown at 130.

The values of the weighting factors $\alpha(Q_T, Q_C)$, $\beta(Q_T, Q_C)$, and $\gamma(Q_T, Q_C)$ are then used to determine how to fuse the scores s_T , s_C , as shown at 132. This fused score can be

25 used to make a decision, as shown at 134, as to whether the first and second voice segments were obtained from the registered user.

Figure 10 is a schematic diagram, illustrating a second embodiment of the authentication process that is generally as described above. Features that are the

30 same as features of the process shown in Figure 9 are indicated by the same reference numerals.

Thus, a first background model 110 for use in a text-constrained (TC) speaker recognition process and a second background model 112 for use in a text-independent

35 (TI) speaker recognition process are taken as starting points. During an initial enrolment, user speech samples are used to form a text-constrained user voiceprint

(that is, the set of coefficients that define the user model) 114 and a text-independent user voiceprint 116.

5 A trigger detection block 118 detects that a trigger phrase has been spoken and a text-constrained speaker recognition or voice biometrics process 120 is performed, based on the first voice segment (namely the user speech from the trigger phrase), the first background model 110, and the text-constrained user voiceprint 114, to produce a score s_T . Also, a text-independent speaker recognition or voice biometrics process 122 is performed, based on the second voice segment (namely the user speech received
10 after the trigger phrase, which is expected to be a command), the second background model 112, and the text-independent user voiceprint 116, to produce a score s_C .

A voice activity detector 140 is used to obtain a measure of the net speech n_C contained in the second voice segment.

15 The functions that map the quality measures Q_T , Q_C to the weighting factors α , β , and γ have been obtained during a development phase and are stored at 128. Thus, the values of the net speech factor n_C is used to obtain values of the weighting factors $\alpha(n_C)$, $\beta(n_C)$, and $\gamma(n_C)$ as shown at 142.

20 The values of the weighting factors $\alpha(n_C)$, $\beta(n_C)$, and $\gamma(n_C)$ are then used to determine how to fuse the scores s_T , s_C , as shown at 132. This fused score can be used to make a decision, as shown at 134 in Figure 10, as to whether the first and second voice segments were obtained from the registered user.

25 The decision illustrated as 134 in Figure 10 is equivalent to the decision shown as 62 in Figure 3. This authentication result from this block 62 can then be used for any desired purpose. For example, the result may be passed to a further decision block 64, to be discussed in more detail below. More generally, if the user is authenticated, a decision
30 may be made to accept that the user is the registered user of the device, as shown at step 74, and this decision may for instance prompt some physical feedback to the user, for example audible, visual or haptic feedback, or some other default action that may occur without further speech input from the user, or allow a response some other input from the user, for instance via a keyboard, to execute some command.

35

In the example shown in Figure 3, if the user is not authenticated based on the combined result, the process passes to step 68. In this example, if the voice biometric input is not considered sufficient to authorise the user, a second authentication factor is requested. For example, the user may be asked to enter a PIN number or password
5 through a keypad of the device, or to provide additional biometric authentication, such as a fingerprint scan.

In step 68, it is determined based on this second authentication factor whether the user is authenticated. This authentication result can then be used for any desired purpose.
10 For example, the result may be passed to a further decision block 70, to be discussed in more detail below. More generally, if the user is authenticated, a decision may be made to accept that the user is the registered user of the device, as shown at step 74 as discussed above.

15 In the example shown in Figure 3, if the user is not authenticated based on the second authentication factor, the process passes to step 72, and the user is rejected. The form that this rejection will take will depend on the action that the user was trying to perform.

The decision to accept or reject the user will therefore have a high degree of reliability,
20 because it takes account of more of the user's speech, compared with a system in which the authentication is performed based only on the trigger phrase.

The example shown in Figure 3 relates to a situation in which the user is attempting to send a speech command to the device, in order to control an aspect of the operation of
25 the device.

Thus, it is expected that the second voice segment, received in step 56, is a command uttered by the user. In step 76, a speech recognition process is performed on the second voice segment. In some embodiments, this may be achieved by transmitting
30 data representing the speech to a speech recognition engine located remotely from the device (for example, in the cloud), using the interface 30 of the device 10. In other embodiments, the speech recognition process is performed in the device 10. The speech recognition process obtains the content and meaning of the second voice segment.

35

In step 78 of this example embodiment, it is determined whether the command is such that it requires authentication. This determination may be made on the basis of a pre-defined list of commands requiring authentication or not requiring authentication, or it may be made on the basis of data or processing modules requested by the command.

5 For example, if the spoken command asks for publicly available information, it may be determined at step 78 that no authentication is required, and hence the process passes to step 80, in which the content of the command may be used to generate a query to an internet search engine in order to be able to supply that information to the user.

10 However, in other cases, for example if the spoken command relates to personal information, it may be determined at step 78 that authentication is required. In that case, the process passes to step 64, in which it is determined whether the user has been authenticated based on the results of the speaker recognition process. If the user has not been authenticated, the process passes to step 70, in which it is determined
15 whether the user has been authenticated, based on the combination of the speaker recognition process and the second authentication factor, as described above.

If the user has not been authenticated, the process passes to step 72, and that command is rejected. If the user has been authenticated, the process passes to step
20 80, and the command is executed.

The decision to execute the command, or not to execute the command, may be communicated to the user, for example by providing audible, visual or haptic feedback to the user.

25

In the example shown in Figure 3, the process of authentication and the process of speech recognition are performed in parallel, in order to save time.

However, in other examples, these processes may be performed sequentially, with the
30 performance of one possibly being conditional on the result of the other.

For example, the process of performing the speech recognition may be performed in all cases, with process of performing speaker recognition being initiated only if it is determined in step 78 that the command is one that requires authentication of the user
35 before that command is performed.

As another example, the speaker recognition process may be performed in all cases, with the speech recognition process being initiated only if it is determined in step 62 (or in step 68) that it is the registered user who is speaking.

- 5 As a further example, the result of the first speaker recognition process may be compared against a threshold (perhaps a low threshold that is relatively easily satisfied), with the speech recognition process of step 76 being performed only if that first speaker recognition process provides a preliminary decision that it is the registered user who is speaking.

10

The examples described above refer to a first voice segment (for example representing a trigger phrase) and one second voice segment (for example representing a command), with the authentication decision being made based on the speaker recognition processes performed on those two voice segments.

15

In further embodiments, for example for use in situations in which there is an extended interaction between the user and the voice user interface, or the user is expected to speak for a longer period of time, the reliability of the process of authentication may be further improved by using more of the user's speech as an input.

20

For example, separate speech inputs, which are provided by the user in response to prompts from the voice user interface following a trigger phrase, might be considered as second, third, fourth etc voice segments. Similarly, a longer period of speech might be divided into second, third, fourth etc voice segments.

25

In such cases, a third speaker recognition process might be performed on the third speech voice segment, a fourth speaker recognition process might be performed on the fourth voice segment, and so on.

- 30 The results of these processes can then be combined. For example a combined result can be obtained by forming a weighted sum of the results from the separate processes. In that case, an interim authentication result can be output after performing the second speaker recognition process, with updated authentication results being output after performing the third and fourth speaker recognition processes, and so on.

35

Alternatively, a result can be obtained based on the second voice segment, with a new result being based on the second and third voice segments considered together, and a further new result being based on the second, third and fourth voice segments all being considered together, and so on, perhaps with older voice segments being discarded as the process continues.

Thus, the process of authenticating the speaker can be performed continuously.

The skilled person will thus recognise that some aspects of the above-described apparatus and methods, for example the calculations performed by the processor may be embodied as processor control code, for example on a non-volatile carrier medium such as a disk, CD- or DVD-ROM, programmed memory such as read only memory (Firmware), or on a data carrier such as an optical or electrical signal carrier. For many applications embodiments of the invention will be implemented on a DSP (Digital Signal Processor), ASIC (Application Specific Integrated Circuit) or FPGA (Field Programmable Gate Array). Thus the code may comprise conventional program code or microcode or, for example code for setting up or controlling an ASIC or FPGA. The code may also comprise code for dynamically configuring re-configurable apparatus such as re-programmable logic gate arrays. Similarly the code may comprise code for a hardware description language such as VerilogTM or VHDL (Very high speed integrated circuit Hardware Description Language). As the skilled person will appreciate, the code may be distributed between a plurality of coupled components in communication with one another. Where appropriate, the embodiments may also be implemented using code running on a field-(re)programmable analogue array or similar device in order to configure analogue hardware

Embodiments of the invention may be arranged as part of an audio processing circuit, for instance an audio circuit which may be provided in a host device. A circuit according to an embodiment of the present invention may be implemented as an integrated circuit.

Embodiments may be implemented in a host device, especially a portable and/or battery powered host device such as a mobile telephone, an audio player, a video player, a PDA, a mobile computing platform such as a laptop computer or tablet and/or a games device for example. Embodiments of the invention may also be implemented wholly or partially in accessories attachable to a host device, for example in active

speakers or headsets or the like. Embodiments may be implemented in other forms of device such as a remote controller device, a toy, a machine such as a robot, a home automation controller or suchlike.

- 5 It should be noted that the above-mentioned embodiments illustrate rather than limit the invention, and that those skilled in the art will be able to design many alternative embodiments without departing from the scope of the appended claims. The word “comprising” does not exclude the presence of elements or steps other than those listed in a claim, “a” or “an” does not exclude a plurality, and a single feature or other
- 10 unit may fulfil the functions of several units recited in the claims. Any reference signs in the claims shall not be construed so as to limit their scope.

CLAIMS

1. A method of processing a received signal representing a user's speech, the method comprising:
 - 5 performing a first speaker recognition process on a first portion of the received signal, to obtain a first output result;
 - performing a second speaker recognition process on a second portion of the received signal that is different from the first portion of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from
 - 10 the first speaker recognition process; and
 - combining the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.
2. A method according to claim 1, comprising applying respective weighting values
- 15 to the first and second output results and combining the weighted results.
3. A method according to claim 2, comprising forming a weighted sum of the first and second output results.
- 20 4. A method according to claim 2 or 3, wherein the weighting values are predetermined weighting values.
5. A method according to claim 2 or 3, wherein the weighting values are variable weighting values.
- 25 6. A method according to claim 5, wherein at least one of the weighting values is determined based on at least one measure relating to a quality of a respective sample.
7. A method according to claim 6, wherein the weighting value is determined as a
- 30 piece-wise constant function of the measure relating to the quality of the respective sample.
8. A method according to claim 6, wherein the weighting value is determined as a continuous function of the measure relating to the quality of the respective sample.

9. A method according to claim 1, comprising applying the first and second output results as inputs to a neural network, and obtaining the combined output result as an output of the neural network.

5 10. A method according to claim 9, further comprising applying at least one measure relating to a quality of a respective sample as an input to the neural network.

11. A method according to claim 9 or 10, wherein the neural network was adapted in a prior development phase.

10

12. A method according to claim 1, comprising combining the first and second output results based on at least one measure relating to a quality of a respective sample.

13. A method according to any one of claims 6, 10 or 12, wherein the at least one
15 measure relating to the quality of the respective sample comprises a measure relating to a property of the respective portion of the received signal.

14. A method according to any one of claims 6, 10 or 12, wherein the at least one
20 measure relating to the quality of the respective sample comprises a measure relating to a property of the speech represented by the respective portion of the received signal.

15. A method according to claim 14, wherein the at least one measure relating to the
25 quality of the respective sample comprises a measure of an amount of speech represented by the respective portion of the received signal.

16. A method according to claim 14, wherein the at least one measure relating to the
30 quality of the respective sample comprises a measure of a degree of fit of the respective portion of the received signal to a background model of the respective speaker recognition process.

17. A method according to claim 13, wherein the at least one measure relating to the
quality of the respective sample comprises a measure relating to a degree of saturation
of the respective portion of the received signal.

35

18. A method according to claim 14, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a signal-to-noise ratio of the respective portion of the received signal.

5 19. A method according to claim 14, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a signal-to-interference ratio of the respective portion of the received signal.

10 20. A method according to claim 14, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a direct-to-reflected ratio of the respective portion of the received signal.

15 21. A method according to claim 14, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a direction from which the respective portion of the received signal was received.

20 22. A method according to claim 14, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a range from which the respective portion of the received signal was received.

23. A method according to any of claims 1 – 22, wherein the first and second speaker recognition processes use different models of the user's speech.

25 24. A method according to any of claims 1 – 23, wherein the first and second speaker recognition processes use different background models.

25. A method according to one of claims 23 and 24, wherein at least one of the models of the user's speech or background models comprises a Hidden Markov Model.

30 26. A method according to one of claims 23 - 25, wherein at least one of the models of the user's speech or background models comprises a Gaussian Mixture Model.

27. A method according to claim 26, wherein the Gaussian Mixture Model uses factor analysis.

28. A method according to claim 26 or 27, wherein the first and second speaker recognition processes use Gaussian Mixture Models with different numbers of Gaussians.

5 29. A method according to any of claims 1 – 23, wherein the first and second speaker recognition process use different verification methods.

30. A method according to any of claims 1 – 29, further comprising performing an antispoofing process on at least one of the first and second portions of the received
10 signal.

31. A method according to claim 30, wherein the antispoofing process comprises computing a likelihood ratio based on a feature vector obtained from the respective portion of the received signal.

15

32. A method according to claim 31, wherein the feature vector comprises one or more metric selected from a spectral ratio, a low frequency ratio, a feature vector squared Mahalanobis distance.

20 33. A method according to claim 30, comprising performing respective antispoofing processes on the first and second portions of the received signal to obtain first and second antispoofing scores respectively.

34. A method according to claim 33, comprising determining that the user is not a
25 registered user if the first or second antispoofing score exceeds a respective threshold.

35. A method according to claim 33, comprising determining that the user is not a registered user if a function of the first and second antispoofing score exceeds a respective threshold.

30

36. A method according to one of claims 30 - 35, and further according to any one of claims 6, 10 or 12, wherein the measure relating to the quality of the respective sample comprises the respective antispoofing score obtained from the respective portion of the received signal.

35

37. A method according to one of claims 33 - 35, comprising combining the first and second output results and the first and second antispoofing scores to obtain the combined output result indicating the likelihood that the user is a registered user.
- 5 38. A method according to claim 37, comprising combining the first and second output results and the first and second antispoofing scores in a classifier.
39. A method according to claim 37, comprising combining the first and second output results and the first and second antispoofing scores in a neural network.
- 10 40. A method according to any one of claims 1 – 39, wherein the first portion of the received signal comprises a trigger phrase and the second portion of the received signal comprises a command.
- 15 41. A method according to any one of claims 1 – 40, wherein the first portion of the received signal corresponds to a first time window and the second portion of the received signal corresponds to a second time window, and wherein the first time window does not overlap the second time window.
- 20 42. A method according to any one of claims 1 – 40, wherein the first portion of the received signal corresponds to a first time window and the second portion of the received signal corresponds to a second time window, and wherein the first time window at least partially overlaps the second time window.
- 25 43. A method according to claim 42, wherein the second time window comprises the first time window.
44. A method according to any one of claims 1 – 43, comprising performing at least one further speaker recognition process on at least one further portion of the received
30 signal to obtain at least one respective further output result; and
combining the at least one further output result with the first and second output results to obtain the combined output result indicating a likelihood that the user is a registered user.
- 35 45. A method according to any one of claims 1 – 44, further comprising performing speech recognition on at least the first portion of the received signal.

46. A method according to any one of claims 1 – 44, further comprising providing feedback to the user based on the combined output result.

5 47. A method according to claim 46, wherein the feedback comprises audible, visual or haptic feedback.

48. A method according to claim 46, comprising allowing or preventing a further action by the user based on the combined output result.

10

49. A device for processing a received signal representing a user's speech, for performing speaker recognition, wherein the device is configured to:

perform a first speaker recognition process on a first portion of the received signal, to obtain a first output result;

15

perform a second speaker recognition process on a second portion of the received signal that is different from the first portion of the received signal, to obtain a second output result, wherein the second speaker recognition process is different from the first speaker recognition process; and

20 combine the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.

50. A device as claimed in claim 49, wherein the device comprises a mobile telephone, an audio player, a video player, a PDA, a mobile computing platform, a games device, a remote controller device, a toy, a machine, or a home automation
25 controller.

51. A device as claimed in claim 49 or 50, further configured for performing speech recognition on at least a portion of the received signal.

30 52. A device as claimed in claim 49, 50 or 51, further configured for transferring at least a portion of the received signal to a remote device for speech recognition.

53. A device as claimed in claim 52, further configured for receiving a result of the speech recognition.

35

54. An integrated circuit device for processing a received signal representing a user's speech, for performing speaker recognition, wherein the integrated circuit device is configured to:

- perform a first speaker recognition process on a first portion of the received
5 signal, to obtain a first output result;
- perform a second speaker recognition process on a second portion of the
received signal that is different from the first portion of the received signal, to obtain a
second output result, wherein the second speaker recognition process is different from
the first speaker recognition process; and
- 10 combine the first and second output results to obtain a combined output result
indicating a likelihood that the user is a registered user.

55. An integrated circuit device as claimed in claim 54, wherein the first and second
speaker recognition processes use at least one user or background model stored in
15 said device.

56. An integrated circuit device as claimed in claim 54 or 55, wherein the first and
second speaker recognition processes use at least one user or background model
stored in a separate integrated circuit device.

20

57. A computer program product, comprising a computer-readable tangible medium,
and instructions for performing a method according to any one of claims 1 to 48.

58. A method of processing a received signal representing a user's speech, the
25 method comprising:

- performing a first speaker recognition process on a first portion of the received
signal, to obtain a first output result;
- performing a second speaker recognition process on a second portion of the
received signal, to obtain a second output result;
- 30 combining the first and second output results to obtain a combined output result
indicating a likelihood that the user is a registered user;
- determining if the indicated likelihood that the user is a registered user is higher
than a threshold value;
- performing a speech recognition process on the second portion of the received
35 signal to obtain a command contained in the second portion of the received signal; and

acting on the obtained command if the indicated likelihood is higher than the threshold value.

59. A method according to claim 58, comprising applying respective weighting values
5 to the first and second output results and combining the weighted results.

60. A method according to claim 59, comprising forming a weighted sum of the first and second output results.

10 61. A method according to claim 59 or 60, wherein the weighting values are predetermined weighting values.

62. A method according to claim 59 or 60, wherein the weighting values are variable
weighting values.

15 63. A method according to claim 62, wherein at least one of the weighting values is determined based on at least one measure relating to a quality of a respective sample.

64. A method according to claim 63, wherein the weighting value is determined as a
20 piece-wise constant function of the measure relating to the quality of the respective sample.

65. A method according to claim 63, wherein the weighting value is determined as a
continuous function of the measure relating to the quality of the respective sample.

25 66. A method according to claim 58, comprising applying the first and second output results as inputs to a neural network, and obtaining the combined output result as an output of the neural network.

30 67. A method according to claim 66, further comprising applying at least one measure relating to a quality of a respective sample as an input to the neural network.

68. A method according to claim 66 or 67, wherein the neural network was adapted
in a prior development phase.

35

69. A method according to claim 58, comprising combining the first and second output results based on at least one measure relating to a quality of a respective sample.

5 70. A method according to any one of claims 63, 67 or 69, wherein the at least one measure relating to the quality of the respective sample comprises a measure relating to a property of the respective portion of the received signal.

71. A method according to any one of claims 63, 67 or 69, wherein the at least one
10 measure relating to the quality of the respective sample comprises a measure relating to a property of the speech represented by the respective portion of the received signal.

72. A method according to claim 71, wherein the at least one measure relating to the
15 quality of the respective sample comprises a measure of an amount of speech represented by the respective portion of the received signal.

73. A method according to claim 71, wherein the at least one measure relating to the
20 quality of the respective sample comprises a measure of a degree of fit of the respective portion of the received signal to a background model of the respective speaker recognition process.

74. A method according to claim 70, wherein the at least one measure relating to the
25 quality of the respective sample comprises a measure relating to a degree of saturation of the respective portion of the received signal.

75. A method according to claim 71, wherein the at least one measure relating to the
30 quality of the respective sample comprises a measure of a signal-to-noise ratio of the respective portion of the received signal.

76. A method according to claim 71, wherein the at least one measure relating to the
quality of the respective sample comprises a measure of a signal-to-interference ratio
of the respective portion of the received signal.

77. A method according to claim 71, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a direct-to-reflected ratio of the respective portion of the received signal.
- 5 78. A method according to claim 71, wherein the at least one measure relating to the quality of the respective sample comprises a measure of a direction from which the respective portion of the received signal was received.
79. A method according to claim 71, wherein the at least one measure relating to the
10 quality of the respective sample comprises a measure of a range from which the respective portion of the received signal was received.
80. A method according to any of claims 58 – 79, wherein the first and second speaker recognition processes are different.
- 15 81. A method according to claim 80, wherein the first and second speaker recognition processes use different models of the user's speech.
82. A method according to claim 80 or 81, wherein the first and second speaker
20 recognition processes use different background models.
83. A method according to one of claims 81 and 82, wherein at least one of the models of the user's speech or background models comprises a Hidden Markov Model.
- 25 84. A method according to one of claims 81 - 83, wherein at least one of the models of the user's speech or background models comprises a Gaussian Mixture Model.
85. A method according to claim 84, wherein the Gaussian Mixture Model uses factor analysis.
30
86. A method according to claim 84 or 85, wherein the first and second speaker recognition processes use Gaussian Mixture Models with different numbers of Gaussians.
- 35 87. A method according to any of claims 80 – 82, wherein the first and second speaker recognition process use different verification methods.

88. A method according to any of claims 58 – 87, further comprising performing an antispoofing process on at least one of the first and second portions of the received signal.

5

89. A method according to claim 88, wherein the antispoofing process comprises computing a likelihood ratio based on a feature vector obtained from the respective portion of the received signal.

10 90. A method according to claim 89, wherein the feature vector comprises one or more metric selected from a spectral ratio, a low frequency ratio, a feature vector squared Mahalanobis distance.

15 91. A method according to claim 90, comprising performing respective antispoofing processes on the first and second portions of the received signal to obtain first and second antispoofing scores respectively.

92. A method according to claim 91, comprising determining that the user is not a registered user if the first or second antispoofing score exceeds a respective threshold.

20

93. A method according to claim 91, comprising determining that the user is not a registered user if a function of the first and second antispoofing score exceeds a respective threshold.

25 94. A method according to one of claims 88 - 93, and further according to any one of claims 58, 62 or 64, wherein the measure relating to the quality of the respective sample comprises the respective antispoofing score obtained from the respective portion of the received signal.

30 95. A method according to one of claims 91 - 93, comprising combining the first and second output results and the first and second antispoofing scores to obtain the combined output result indicating the likelihood that the user is a registered user.

35 96. A method according to claim 95, comprising combining the first and second output results and the first and second antispoofing scores in a classifier.

97. A method according to claim 95, comprising combining the first and second output results and the first and second antispoofing scores in a neural network.

98. A method according to any one of claims 58 – 97, wherein the first portion of the received signal comprises a trigger phrase and the second portion of the received signal comprises a command.

99. A method according to any one of claims 58 – 98, wherein the first portion of the received signal corresponds to a first time window and the second portion of the received signal corresponds to a second time window, and wherein the first time window does not overlap the second time window.

100. A method according to any one of claims 58 – 98, wherein the first portion of the received signal corresponds to a first time window and the second portion of the received signal corresponds to a second time window, and wherein the first time window at least partially overlaps the second time window.

101. A method according to claim 100, wherein the second time window comprises the first time window.

102. A method according to any one of claims 58 – 101, comprising performing at least one further speaker recognition process on at least one further portion of the received signal to obtain at least one respective further output result; and combining the at least one further output result with the first and second output results to obtain the combined output result indicating a likelihood that the user is a registered user.

103. A method according to any one of claims 58 – 102, further comprising performing speech recognition on at least the first portion of the received signal.

104. A method according to any one of claims 58 – 103, further comprising providing feedback to the user based on the combined output result.

105. A method according to claim 104, wherein the feedback comprises audible, visual or haptic feedback.

106. A method according to claim 104, comprising allowing or preventing a further action by the user based on the combined output result.

107. A device for processing a received signal representing a user's speech, for
5 performing speaker recognition, wherein the device is configured to:

perform a first speaker recognition process on a first portion of the received signal, to obtain a first output result;

perform a second speaker recognition process on a second portion of the received signal that is different from the first portion of the received signal, to obtain a
10 second output result, wherein the second speaker recognition process is different from the first speaker recognition process; and

combine the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.

15 108. A device as claimed in claim 107, wherein the device comprises a mobile telephone, an audio player, a video player, a PDA, a mobile computing platform, a games device, a remote controller device, a toy, a machine, or a home automation controller.

20 109. A device as claimed in claim 107 or 108, further configured for performing speech recognition on at least a portion of the received signal.

110. A device as claimed in claim 107, 108 or 109, further configured for transferring at least a portion of the received signal to a remote device for speech recognition.

25 111. A device as claimed in claim 110, further configured for receiving a result of the speech recognition.

112. An integrated circuit device for processing a received signal representing a user's
30 speech, for performing speaker recognition, wherein the integrated circuit device is configured to:

perform a first speaker recognition process on a first portion of the received signal, to obtain a first output result;

perform a second speaker recognition process on a second portion of the
35 received signal that is different from the first portion of the received signal, to obtain a

second output result, wherein the second speaker recognition process is different from the first speaker recognition process; and

combine the first and second output results to obtain a combined output result indicating a likelihood that the user is a registered user.

5

113. An integrated circuit device as claimed in claim 112, wherein the first and second speaker recognition processes use at least one user or background model stored in said device.

10 114. An integrated circuit device as claimed in claim 112 or 113, wherein the first and second speaker recognition processes use at least one user or background model stored in a separate integrated circuit device.

15 115. A computer program product, comprising a computer-readable tangible medium, and instructions for performing a method according to any one of claims 58 to 106.

116. A non-transitory computer readable storage medium having computer-executable instructions stored thereon that, when executed by processor circuitry, cause the processor circuitry to perform a method according to any one of claims 1 to 48.

20

117. A non-transitory computer readable storage medium having computer-executable instructions stored thereon that, when executed by processor circuitry, cause the processor circuitry to perform a method according to any one of claims 58 to 106.

25




Figure 1




Figure 2

3/10




Figure 3




Figure 4




Figure 5




Figure 6




Figure 7




Figure 8




Figure 9




Figure 10

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2016/051954

A. CLASSIFICATION OF SUBJECT MATTER

INV. G10L17/10 G10L17/00
 ADD. G10L15/22

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G10L

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

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

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2010/131273 A1 (ALEY-RAZ ALMOG [IL] ET AL) 27 May 2010 (2010-05-27)	1-18, 23-75, 80-117
A	paragraph [0044] - paragraph [0045] paragraph [0047] paragraph [0049] paragraph [0052] - paragraph [0060] paragraph [0071] paragraph [0094] figure 1	19-22, 76-79
A	----- JOSEPH P CAMPBELL ET AL: "Fusing High- and Low-Level Features for Speaker Recognition", 20030901, 4 September 2003 (2003-09-04), page 2665, XP007007102, section 4.3 ----- -/-	9-11, 66-68



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

21 March 2017

Date of mailing of the international search report

07/04/2017

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel. (+31-70) 340-2040,
 Fax: (+31-70) 340-3016

Authorized officer

Schneider, Daniel

INTERNATIONAL SEARCH REPORT

International application No

PCT/GB2016/051954

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	WO 2008/047339 A2 (NICE SYSTEMS LTD [IL]; ZIGEL YANIV [IL]; WASSERBLAT MOSHE [IL]) 24 April 2008 (2008-04-24) page 17, line 29 - page 18, line 15 -----	15,17, 18,72, 74,75
A	EP 2 797 078 A1 (AGNITIO S L [ES]) 29 October 2014 (2014-10-29) paragraph [0062] - paragraph [0063] -----	16,18, 73,75
A	US 2014/257812 A1 (MOZER TODD F [US]) 11 September 2014 (2014-09-11) paragraph [0022] -----	41,42, 99,100
A	US 2014/249817 A1 (HART GREGORY MICHAEL [US] ET AL) 4 September 2014 (2014-09-04) paragraph [0071] -----	58-106, 115,117

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2016/051954

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2010131273 A1	27-05-2010	EP 2361430 A2	31-08-2011
		US 2010131273 A1	27-05-2010
		US 2013226582 A1	29-08-2013
		US 2015046162 A1	12-02-2015
		WO 2010061344 A2	03-06-2010
WO 2008047339 A2	24-04-2008	US 2008195387 A1	14-08-2008
		WO 2008047339 A2	24-04-2008
EP 2797078 A1	29-10-2014	DK 2797078 T3	23-01-2017
		EP 2797078 A1	29-10-2014
		ES 2600461 T3	09-02-2017
		US 2016307572 A1	20-10-2016
		WO 2014173919 A1	30-10-2014
US 2014257812 A1	11-09-2014	CN 103827963 A	28-05-2014
		CN 105719647 A	29-06-2016
		US 2013080167 A1	28-03-2013
		US 2014257812 A1	11-09-2014
		WO 2013048876 A1	04-04-2013
US 2014249817 A1	04-09-2014	EP 2965314 A1	13-01-2016
		US 2014249817 A1	04-09-2014
		WO 2014137751 A1	12-09-2014

INTERNATIONAL SEARCH REPORT

International application No.
PCT/GB2016/051954

Box No. II Observations where certain claims were found unsearchable (Continuation of item 2 of first sheet)

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

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

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

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

see additional sheet

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

Remark on Protest

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

FURTHER INFORMATION CONTINUED FROM PCT/ISA/ 210

This International Searching Authority found multiple (groups of) inventions in this international application, as follows:

1. claims: 1-57, 107-114, 116

merging results from different speaker recognition systems

2. claims: 58-106, 115, 117

perform speech recognition on the portions and execute
detected command
