

(12) STANDARD PATENT
(19) AUSTRALIAN PATENT OFFICE

(11) Application No. **AU 2016333816 B2**

(54) Title
Assessment of a pulmonary condition by speech analysis

(51) International Patent Classification(s)
A61B 5/00 (2006.01)

(21) Application No: **2016333816** (22) Date of Filing: **2016.10.05**

(87) WIPO No: **WO17/060828**

(30) Priority Data

(31) Number	(32) Date	(33) Country
62/238,732	2015.10.08	US

(43) Publication Date: **2017.04.13**

(44) Accepted Journal Date: **2018.09.27**

(71) Applicant(s)
Cordio Medical Ltd.

(72) Inventor(s)
Lotan, Chaim;Kremer-Tal, Sigal;Lotan, Aviv;Schlik, Zeev;Gemer, Avinoam;Snir, Yehuda;Sasson, Yonatan;Sheinkerman, Margarita

(74) Agent / Attorney
FPA Patent Attorneys Pty Ltd, ANZ Tower 161 Castlereagh Street, Sydney, NSW, 2000, AU

(56) Related Art
US 20120265024 A1



(51) International Patent Classification:
A61B 5/00 (2006.01)

(21) International Application Number:
PCT/IB2016/055954

(22) International Filing Date:
5 October 2016 (05.10.2016)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
62/238,732 8 October 2015 (08.10.2015) US

(71) Applicant: CORDIO MEDICAL LTD. [IL/IL]; 6 Yehonatan Netanyahu Street, 6037604 Or-Yehuda (IL).

(72) Inventors: LOTAN, Chaim; 12 Shachar Street, Beit Hakerem, 96263 Jerusalem (IL). KREMER-TAL, Sigal; 59 Hameyasdim Street, 47245 Ramat Hasharon (IL). LOTAN, Aviv; 22 Refidim Street, 6998243 Tel Aviv (IL). SCHLIK, Zeev; 11 Hacabaim Street, 5225511 Ramat Gan (IL). GEMER, Avinoam; 1 Tzahal Street, 5545101 Kiryat Ono (IL). SNIR, Yehuda; 12A Hanarkisim Street, 36073 Kiryat Tivon (IL). SASSON, Yonatan; 32 Groneman

Street, 6997235 Tel Aviv (IL). SHEINKERMAN, Margarita; 29 Zabotinski Street, 6228705 Tel Aviv (IL).


(74) Agent: D. KLIGLER I. P. SERVICES LTD.; P. O Box 57651, 61576 Tel Aviv (IL).

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

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK,

[Continued on next page]

(54) Title: ASSESSMENT OF A PULMONARY CONDITION BY SPEECH ANALYSIS



(57) Abstract: Described embodiments include apparatus that includes a network interface (28) and a processor (30). The processor is configured to receive, via the network interface, speech of a subject (22) who suffers from a pulmonary condition related to accumulation of excess fluid, to identify, by analyzing the speech, one or more speech-related parameters of the speech, to assess, in response to the speech-related parameters, a status of the pulmonary condition, and to generate, in response thereto, an output indicative of the status of the pulmonary condition. Other embodiments are also described.





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

— before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

Published:

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

ASSESSMENT OF A PULMONARY CONDITION BY SPEECH ANALYSIS

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims the benefit of US Provisional Application No. 62/238,732, entitled "Assessment of a pulmonary condition using voice analysis," filed October 8, 2015,
5 whose disclosure is incorporated herein by reference.

FIELD OF THE INVENTION

Embodiments of the present invention relate to the diagnosis and treatment of medical conditions, particularly pulmonary conditions.

BACKGROUND

10 Pulmonary edema is an accumulation of fluid within the parenchyma and air spaces of the lungs. It leads to impaired gas exchange and may cause respiratory failure.

Pleural effusion is excess fluid that accumulates in the pleural cavity, the fluid-filled space that surrounds the lungs. This excess fluid can impair breathing by limiting the expansion of the lungs.

15 US Patent Application Publication 2015/0073306, whose disclosure is incorporated herein by reference, describes a method of operating a computational device to process patient sounds, the method comprising the steps of: extracting features from segments of said patient sounds; and classifying the segments as cough or non-cough sounds based upon the extracted features and predetermined criteria; and presenting a diagnosis of a disease related state on a display under control
20 of the computational device based on segments of the patient sounds classified as cough sounds.

US Patent 7,267,652, whose disclosure is incorporated herein by reference, describes systems and methods for processing respiratory signals derived generally from respiratory plethysmography, and especially from respiratory inductive plethysmographic sensors mounted on a garment for ambulatory recording. The systems and methods provide signal filtering for artifact rejection, and
25 calibration of sensor data to produce outputs indicative of lung volumes. This reference also describes systems and methods directed to processing lung volume signals, however measured or derived, to provide determination of respiratory parameters and recognition of selected respiratory events.

US Patent Application Publication 2011/0184250, whose disclosure is incorporated herein by
30 reference, describes a computer-implemented method and system for assisting a plurality of patients

manage chronic health conditions. The method, for each patient, comprises: (a) receiving information from the patient or a member of a patient care network on an expected patient activity at a given future time period; (b) determining expected transient local ambient conditions in the patient's surroundings during the expected patient activity at the given future time period; (c) predicting health exacerbations for the patient using a stored computer model of the patient based on a desired patient control set-point range, the expected patient activity, and the expected transient local ambient conditions; and (d) proactively sending a message to the patient or a member of the patient care network before the given future time period, the message alerting the patient or a member of the patient care network of the predicted health exacerbations for the patient and identifying one or more corrective actions for the patient to avoid or mitigate the predicted health exacerbations.

PCT Publication WO 2014/037843, whose disclosure is incorporated herein by reference, describes a computerized method and system for measuring a user's lung capacity and stamina, to detect Chronic Heart Failure, COPD or Asthma, comprising: providing a client application on a user's mobile communication device, said client application comprising executable computer code for: instructing the user to fill his lungs with air and utter vocal sounds within a certain range of loudness (decibels) while exhaling; receiving and registering by the mobile communication device said user's vocal sounds; stopping the registering of vocal sounds; measuring the length of the vocal sounds receiving time within said range of loudness, said time proportionate the user's lung volumes; and displaying the length of sound received time results on the mobile communication device screen.

Reference to any prior art in the specification is not an acknowledgment or suggestion that this prior art forms part of the common general knowledge in any jurisdiction or that this prior art could reasonably be expected to be understood, regarded as relevant, and/or combined with other pieces of prior art by a skilled person in the art.

SUMMARY OF THE INVENTION

As used herein, except where the context requires otherwise, the term "comprise" and variations of the term, such as "comprising", "comprises" and "comprised", are not intended to exclude further features, components, integers or steps.

There is provided, in accordance with some embodiments of the present invention, apparatus comprising a network interface and a processor. The processor is configured to receive, via the network interface, speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid, to identify, by analyzing the speech, one or more speech-related parameters of the speech, to assess, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and to

generate, in response thereto, an output indicative of the status of the pulmonary condition.

In some embodiments, the processor is configured to analyze the speech by performing a spectral analysis of the speech.

5 In some embodiments, the processor is configured to analyze the speech by performing a cepstral analysis of the speech.

In some embodiments, the processor is further configured to identify, by analyzing the speech, a meaning of the speech, and the processor is configured to assess the status in response to the meaning.

10 In some embodiments, the processor is further configured to prompt the subject to provide, by orally responding to a question, the speech.

In some embodiments, the processor is configured to prompt the subject to provide the speech by:

placing a call to the subject, and

upon the subject answering the call, asking the question to the subject.

15 In some embodiments, the condition is selected from the group consisting of: pulmonary edema, and pleural effusion.

In some embodiments, the processor is configured to generate the output by providing an alert indicative of a deterioration in the status.

20 In some embodiments, the speech-related parameters include at least one formant frequency of the speech.

In some embodiments, the processor is configured to assess the status of the pulmonary condition by comparing the formant frequency to a baseline frequency.

In some embodiments, the processor is configured to analyze the speech by:

for each frame of a time-domain signal that represents the speech:

25 computing a cepstrum of the frame,

subsequently, applying a low-pass cepstral lifter to the cepstrum of the frame,

subsequently, inverse-transforming the liftered cepstrum to a frequency domain,

and

30 subsequently, identifying a formant frequency of the frame, by identifying a peak of the inverse-transformed liftered cepstrum, and

identifying the formant frequency, based on the respective formant frequencies of the

frames.

In some embodiments, the processor is configured to analyze the speech by:

selecting at least one portion of the speech for analysis, based on one or more metrics selected from the group consisting of: an energy level of the speech, a stability of a pitch of the speech, a confidence in a calculation of the pitch of the speech, and a stability of a formant frequency of the speech, and

subsequently, analyzing the selected portion of the speech.

In some embodiments, the processor is configured to assess the status of the pulmonary condition by applying a machine-learned model to the speech-related parameters.

In some embodiments, the speech-related parameters include a set of mel-frequency cepstrum coefficients (MFCCs).

In some embodiments, the speech-related parameters include a set of derivatives of mel-frequency cepstrum coefficients (MFCCs).

In some embodiments, the speech-related parameters include a set of relative spectral transform - perceptual linear prediction (RASTA-PLP) cepstrum components.

In some embodiments, the speech-related parameters include a set of relative spectral transform - perceptual linear prediction (RASTA-PLP) spectrum components.

In some embodiments, the processor is configured to assess the status of the pulmonary condition without identifying any wheezing sounds, in the speech.

In some embodiments, the processor is configured to analyze the speech of the subject without first prompting the subject to provide the speech.

There is further provided, in accordance with some embodiments of the present invention, a method that includes, using a processor, by analyzing speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid, identifying one or more speech-related parameters of the speech. The method further includes, in response to the speech-related parameters, without having identified any cough sounds in the speech, automatically assessing a status of the pulmonary condition, and, in response to the assesment, generating an output indicative of the status of the pulmonary condition.

There is further provided, in accordance with some embodiments of the present invention, a system that includes a mobile communication terminal, comprising a first processor, and a server. The server includes a second processor, the first processor and the second

processor being configured to cooperatively carry out a process that includes: receiving speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid, identifying, by analyzing the speech, one or more speech-related parameters of the speech, assessing, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and generating, in response thereto, an output indicative of the status of the pulmonary condition.

There is further provided, in accordance with some embodiments of the present invention, a computer software product including a tangible non-transitory computer-readable medium in which program instructions are stored. The instructions, when read by a processor, cause the processor to receive speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid, to identify, by analyzing the speech, one or more speech-related parameters of the speech, to assess, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and to generate, in response to the assessment, an output indicative of the status of the pulmonary condition.

There is further provided, in accordance with some embodiments of the present invention, apparatus that includes an acoustic sensor and a processor. The processor is configured to receive, via the acoustic sensor, speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid, to identify, by analyzing the speech, one or more speech-related parameters of the speech, to assess, in response to the speech-related parameters, a status of the pulmonary condition without having identified any cough sounds in the speech, and to generate, in response to the assessment, an output indicative of the status of the pulmonary condition.

The present invention will be more fully understood from the following detailed description of embodiments thereof, taken together with the drawings, in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a schematic illustration of a system for assessing the status of a pulmonary condition, in accordance with some embodiments of the present invention;

Fig. 2 is a schematic illustration of a preprocessing method performed by a processor, in accordance with some embodiments of the present invention; and

Fig. 3 is a schematic illustration of a method for analyzing speech of a subject, in accordance with some embodiments of the present invention.

DETAILED DESCRIPTION OF EMBODIMENTS

OVERVIEW

Embodiments of the present invention provide methods and systems that assess the status of a pulmonary condition, such as pulmonary edema or pleural effusion condition, that is related to accumulation of excess fluid in or near the lungs. In some embodiments, a system that assesses the status of a pulmonary condition comprises a smartphone and a remote server. The smartphone records the subject as the subject speaks, and then sends the recorded speech signal to the remote server for analysis. By automatically analyzing the signal, the remote server identifies one or more speech-related parameters of the speech, and, in response to the parameters, assesses the status of the condition. The server then generates an output to the smartphone that drives the smartphone, in turn, to generate an output, such as a visual output,

that indicates the status of the condition. Alternatively or additionally, in the event of a deterioration in the subject's condition, the server may alert an emergency call center, and/or the subject's physician or caregiver, of the deterioration, as described below.

Typically, the speech-related parameters, upon which the assessment is based, include
5 spectral and/or cepstral properties of the speech, such as at least one formant frequency of the speech. Although formant frequencies are usually associated with the vocal tract or upper airway, embodiments of the present invention capitalize on the observation that a change in a formant frequency may also indicate a change in the amount of fluid present in the parenchyma, pleural cavity, or other portion of anatomy that belongs to, or is near, the lungs. In particular, it has
10 been observed that an increase in a formant frequency indicates an increase in the amount of fluid, and *vice versa*. The system may therefore identify a deterioration in the status of the condition, based on an increase in a formant frequency, and *vice versa*.

Alternatively or additionally, the system may assess the status of the condition, by applying a machine-learned model to particular parameters, or "features," extracted from the speech
15 signal. Such a model may include a regressor, which outputs a number indicative of the status. For example, the regressor may output a number that indicates, relative to a baseline, an estimated amount of fluid inside, or near, the lungs. Alternatively or additionally, such a model may include a classifier, which classifies the current status of the pulmonary condition, and/or classifies the current status of the subject based on the current status of the pulmonary condition.
20 For example, such a classifier may classify the subject as "in need of dialysis" or "not in need of dialysis."

It is emphasized that embodiments of the present invention do not necessarily require the identification of cough sounds, wheezing sounds, breathing sounds (e.g., cackles), or any other non-speech-related parameters, in order to assess the status of the pulmonary condition. Rather,
25 the assessment may be based only on speech-related parameters, such as formant frequencies of the subject's speech, and/or any of the other speech-related parameters described herein. Moreover, the subject need not necessarily be prompted, at any particular time, to utter any particular sounds, to say any particular words, or even to speak naturally. Rather, the system may analyze spoken sentences obtained from the subject's normal, daily conversations on his
30 smartphone. (Notwithstanding the above, in some embodiments, the subject may be prompted to utter particular sounds, and/or say particular words, particularly if the analysis of the subject's "natural" speech is inconclusive.)

In general, the embodiments described herein facilitate early detection of, and

intervention with respect to, a deterioration in the subject's condition, such that hospitalization of the subject may not be necessary. For example, the embodiments described herein facilitate evaluating the effectiveness of the particular type and/or dosage of medication that the subject is taking, such that, if a deterioration is observed, the particular type and/or dosage of medication
5 may be adjusted, before hospitalization is required.

SYSTEM DESCRIPTION

Reference is initially made to Fig. 1, which is a schematic illustration of a system 20 for assessing the status of a pulmonary condition, in accordance with some embodiments of the present invention. System 20 comprises a device for communication and/or computing, such as
10 a smartphone 24 or other mobile communication terminal, and a server 26. A subject 22, who suffers from a pulmonary condition such as pulmonary edema or pleural effusion, speaks into smartphone 24, typically during the course of the subject's normal day-to-day conversational activities. The smartphone transmits the subject's speech to server 26, which then analyzes the speech, as further described below, such as to assess the status of the pulmonary condition.
15 (Such an analysis may be performed offline, or in real-time, as the subject is speaking.) The server then returns the results of the analysis to the smartphone, which then generates an output, such as an audio and/or visual output, indicative of the status of the pulmonary condition. For example, in response to the server identifying a deterioration in the status, the server may, via the smartphone, provide an alert indicative of the deterioration, and/or instruct the subject to speak
20 with his physician, take his prescribed medication, and/or adjust the type of, or increase the dosage of, the medication taken. Conversely, in response to an improvement in, or stabilization of, the subject's status, the system may recommend a decrease in a medication dosage.

In some embodiments, the server is configured to communicate outputs, such as the results of an analysis and/or an alert, to a remote terminal, alternatively or additionally to
25 communicating these outputs to the subject's smartphone. Thus, for example, the server may communicate outputs to the subject's physician, caregiver, insurance company, or employer, or to a hospital, an emergency call center, or a data center. Alternatively or additionally, the smartphone may communicate outputs to such remote terminals.

In some embodiments, the smartphone is configured to control a medicine pump. In such
30 embodiments, in response to a deterioration in the subject's condition, the smartphone may automatically increase a dosage of medication delivered to the subject by the pump, without necessarily generating an alert. Subsequently, a follow-up analysis may be performed, to ascertain if the increase in dosage was successful. If the increase in dosage was not successful,

an alert may be generated. Alternatively or additionally, the smartphone may be configured to control a diuretic pump in response to the results of an analysis; for example, the smartphone may automatically change the operating frequency of the pump.

In general, the response to a deterioration typically depends on the magnitude of the deterioration. For example, for a relatively small deterioration, the system may recommend, or automatically implement, an increase in the frequency at which the subject's speech is analyzed, and/or expand the analysis to identify additional parameters that may be relevant to the subject's condition, without necessarily generating an alert. For a larger deterioration, an alert may be generated, and/or the system may recommend, or automatically implement, an increased dosage of medication, such as a diuretic. Alternatively or additionally, the system may recommend that the subject see a doctor. In yet more extreme cases, the system may automatically contact an emergency call center.

Typically, server 26 is remote from subject 22, residing, for example, "in the cloud." The server comprises a network interface, such as a network interface controller (NIC) 28, which receives input from, and transmits output to, the smartphone. The server further comprises a processor 30, which, based on inputs received from the smartphone via NIC 28, computes relevant outputs, and returns these outputs, via NIC 28, to the smartphone. Analogously, the smartphone comprises a processor 34, and a network interface, such as a wireless transceiver (TCVR) 32, which is used by processor 34 to exchange communication with the server. The smartphone further comprises an acoustic sensor, configured to receive sounds that are emitted by the subject.

In alternative embodiments, a local processor, such as processor 34, or the processor of another local computing device, such as a laptop or desktop computer, processes the speech of the subject. For example, the subject may speak into a microphone belonging to a laptop computer, and the processor of the laptop computer may then process the resulting speech signals from the microphone.

Typically, processor 34 performs the relevant functions described herein by executing program instructions provided by a dedicated software application. The software application is typically run continuously, by the processor, in the background, such that the subject may continue regular usage of the smartphone while the application is running. Upon the subject beginning to speak into the phone, the application - without the subject necessarily noticing - begins to record, and/or transmit, the subject's speech. Alternatively, the application may be manually activated by the subject, and/or may be automatically activated by the processor, upon

the subject initiating or answering a call.

Typically, the application instructs processor 34 to send a speech signal to the server for analysis, whenever a particular duration has passed from the most recent analysis, or whenever the server or the subject requests an analysis. Processor 34 may then wait for the subject to engage in a phone call, at which point, following receipt of a speech signal, the processor may send the speech signal to the server. Alternatively, the processor may prompt the subject (e.g., by displaying a message on the smartphone display) to speak into the smartphone, such that a speech signal may be obtained. In some embodiments, the server automatically places a call to the subject (by automatically calling the smartphone), and then, upon the subject answering the call, prompts the subject to speak. For example, the server may ask a question to the subject, such as "Did you take your medicine today?" (The question may be recorded or voice-synthesized.) Upon receiving the subject's answer, the server may analyze the answer for meaning and/or vocal properties, as described below.

Alternatively, the subject may be prompted to speak in any other suitable way, such as by receiving instructions to do so from medical personnel. In some embodiments, the subject may be prompted, e.g., by the smartphone, to perform a particular action prior to speaking, such as taking a certain number of steps, or climbing a certain number of stairs.

Typically, as noted above, processor 30 analyzes natural, "free" speech of the subject. In some cases, however, the subject may be prompted, e.g., by processor 34, to recite particular words or sentences that have been found to facilitate the analysis described herein, and these words or sentences may then be communicated to the server for analysis. (The same particular words or sentences may be used each time the analysis is performed, in order to facilitate a more effective assessment of the subject's condition.)

In some embodiments, the meaning (i.e., semantic content) of the subject's speech is automatically identified, and the assessment of the status of the condition is performed in response to the meaning. For example, the subject may be prompted, e.g., by the smartphone, to provide the speech by orally responding to a particular question. The subject's oral response to the question is then analyzed both for the identification of the speech-related parameters described herein - which describe the vocal properties, such as spectral and/or cepstral properties, of the response - and for the meaning of the response. (The question may be posed visually, on the display of the smartphone, and/or auditorily.) Examples of such questions include "Did you take your medicine today?," "How do you feel?," and "How active were you today?" By identifying both the vocal properties of, and meaning of, the subject's answer, the

status of the subject's condition may be better assessed. (In such embodiments, the subject may be required to select an answer from a predetermined group of candidate answers, to facilitate the analysis.) For example, if the subject responds "Yes" to the question "Did you take your medicine today?," but the speech-related parameters indicate a deterioration in the subject's condition, the processor may ascertain that the dosage of medicine prescribed to the subject is not high enough.

In some cases, processor 30 may analyze particular, possibly meaningless, sounds, such as a prolonged "ahh." (Processor 34 of the smartphone may be configured to prompt the subject to utter such sounds, upon receiving a request from the server to do so.) In analyzing these sounds, processor 30 may identify any one or more of the relevant parameters described herein. These parameters may then be used, optionally in combination with one or more other parameters derived from natural speech of the subject, to assess the status of the subject's condition.

In general, processor 30 may analyze the subject's speech at any suitable predetermined times, separated by any suitable intervals. For example, the processor may analyze the subject's speech immediately after hospitalization (when the subject's condition is relatively stable), one week after hospitalization, and three weeks after hospitalization. Alternatively, for example, the processor may analyze the subject's speech daily (beginning, for example, upon the subject's release from the hospital), or every other day. In some embodiments, the analysis schedule is customized to the subject's perception of the subject's condition; for example, an analysis may be performed upon the subject perceiving an improvement or deterioration in the subject's condition.

In general, each of processor 30 and processor 34 may be embodied as a single processor, or as a cooperatively networked or clustered set of processors. Each of processor 30 and processor 34 is typically a programmed digital computing device comprising a central processing unit (CPU), random access memory (RAM), non-volatile secondary storage, such as a hard drive or CD ROM drive, network interfaces, and/or peripheral devices. Program code, including software programs, and/or data, are loaded into the RAM for execution and processing by the CPU and results are generated for display, output, transmittal, or storage, as is known in the art. The program code and/or data may be downloaded to the processor in electronic form, over a network, for example, or it may, alternatively or additionally, be provided and/or stored on non-transitory tangible media, such as magnetic, optical, or electronic memory. Such program code and/or data, when provided to the processor, produce a machine or special-purpose computer, configured to perform the tasks described herein.

Reference is now made to Fig. 2, which is a schematic illustration of a preprocessing method performed by processor 30, in accordance with some embodiments of the present invention.

In some embodiments, prior to analyzing the speech signal received from the smartphone, processor 30 preprocesses the speech signal, in order to identify those portions of the signal that are most suitable for analysis. Fig. 2 shows, on the left, the original speech signal 36, in the time domain, received from the smartphone. Typically, the processor first separates signal 36 into a plurality of (possibly overlapping) frames, each of which includes several samples. (For example, assuming a sampling rate of 8 kHz, a frame of 40 ms would include 320 samples.) Then, from signal 36, the processor derives one or more metrics, which are used by the processor to select those portions of the signal that are most suitable for analysis.

For example, the processor may compute the energy level 38 of the speech signal, expressed, for example, on a logarithmic scale relative to the maximum energy level in the signal. The processor may then select, for analysis, the largest contiguous portion of the signal whose energy level is above a particular threshold 40. Alternatively or additionally, the processor may compute the pitch 42 of the speech signal, and a corresponding confidence level 44, which indicates the confidence with which the pitch calculation was performed. The processor may then select portions of the signal for analysis, based on the pitch values, and/or the pitch confidence levels. For example, the processor may select the largest contiguous portion of the signal in which the stability of the pitch is above a particular threshold, i.e., the pitch changes by less than a particular threshold, and/or the confidence of the pitch calculation is above a particular threshold. Alternatively or additionally, the processor may calculate at least one formant frequency 46 of each frame belonging to the signal, and base the selection on the stability of the formant frequency.

Typically, all of the above-described selection criteria are implemented. That is, typically, the processor selects only those portions of the signal whose (i) energy level is above the threshold, (ii) pitch is sufficiently stable, (iii) pitch confidence is sufficiently high, and (iv) formant frequency is sufficiently stable. In particular, following the selection of a portion of a signal based on the energy-level criterion, the processor typically divides the signal into a plurality of larger segments, each segment including several frames, and then selects, for analysis, a consecutive sequence of segments that satisfy the pitch-based and formant-based selection criteria described above. For example, assuming that frames 131 through 680 have sufficient energy (as roughly indicated by Fig. 3), the processor may divide frames 131 through 680 into ten segments S1, S2, ... S10, each segment including 55 frames. Then, a consecutive

subset of these segments (e.g., the second through ninth of these segments) may be selected, based on the pitch-based and formant-based selection criteria.

It is emphasized that the above-described preprocessing typically does not involve cleaning the speech signal in any way. In other words, processor 30 typically does not alter the speech signal, in order to render the signal more suitable for analysis. Rather, as described above, the processor selects any suitable portions of the signal, while preserving the original content of these portions. If an insufficiently-large portion of the signal is found to be suitable for analysis, processor 30 may instruct the smartphone to transmit another signal to the server. Following the above-described preprocessing technique, processor 30 analyzes the portion of the speech signal that was ascertained to be suitable for analysis, such as to identify one or more speech-related parameters of the speech. Then, in response to the parameters, the processor assesses the status of the pulmonary condition. For example, the processor may identify at least one formant frequency of the speech (e.g., the formant frequency F1, F2, F3, and/or F4), and assesses the status of the pulmonary condition in response to the formant frequency. Typically, the processor assesses the status of the pulmonary condition by comparing the formant frequency to a baseline (or "reference") frequency. Such a baseline frequency may be a formant frequency, of the subject's speech, that was measured at an earlier date. If the current formant frequency is greater than this baseline, the subject's condition has deteriorated; conversely, if the current formant frequency is less than this baseline, the subject's condition has improved. If the current formant frequency is approximately equal to the baseline, the subject's condition is stable.

For example, the processor may use, as the baseline, a formant frequency that was measured (e.g., by the processor, using techniques described herein) upon the subject's release from hospital. If the current formant frequency is greater than this baseline, the subject's condition has deteriorated; on the other hand, little change in the formant frequency may indicate that the subject's condition remains stable.

Alternatively, such a baseline frequency may be a target (or "plateau") frequency at which the subject's condition is deemed to be stable, or a "borderline" frequency at which intervention and/or hospitalization is required. (Such baselines may be established separately for each subject, based on the subject's medical history.) The processor may assess the current status of the pulmonary condition, by comparing the current formant frequency to such a baseline. In some embodiments, an alert is generated if the difference between the current formant frequency and the target frequency is greater than a particular threshold, or if the current formant frequency is greater than the borderline frequency.

In some embodiments, the subject's formant frequency is measured over a period of time during which the subject is deemed (e.g., by a physician) to be healthy. The mean of these measurements is then taken as the baseline, and the standard deviation of these measurements is used to define a threshold at which alerts are generated. For example, an alert may be generated
5 if a subsequent formant frequency exceeds the baseline by more than two standard deviations. As a purely illustrative example, the formant frequency F1 may have a baseline of 230 Hz with a standard deviation of 17 Hz, such that an alert is generated if F1 exceeds 264 Hz.

In some cases, such as where the subject's history is not available, the baseline may be based on population averages. In this regard, relevant characteristics of the subject - such as the
10 subject's gender, weight, height, age, or level of physical activity - may be obtained, and used to choose a suitable baseline. For example, a population of individuals may be clustered into a plurality of clusters, based on relevant characteristics of the individuals, and a baseline for the subject may then be selected, by computing the average of the relevant parameter (e.g., the formant frequency) for the cluster in which the subject falls.

In some embodiments, the processor identifies a trend in the formant frequency, and
15 generates an output, such as an alert, based on the identified trend. For example, if the processor identifies an increase in the formant frequency over a certain number of days, the processor may identify a deterioration in the subject's condition, and hence, generate an alert. (Such an alert may be generated even if the formant frequency has not exceeded the threshold that is based on
20 the baseline frequency.) Alternatively, if the processor identifies a decrease in the formant frequency over a certain number of days, the processor may generate an output that indicates an improvement in the subject's condition.

In some embodiments, to identify the formant frequency, the processor first identifies the
25 respective formant frequency of each of the frames in the selected segments. (The processor may first apply a Hamming window to each of the frames, and/or filter each frame with a pre-emphasis filter.) Next, the formant frequency of the speech (i.e., the speech signal as a whole) is identified, based on the respective formant frequencies of the frames. For example, the processor may compute the mean, or median, of the respective formant frequencies of the frames, and identify this metric as the formant frequency of the speech.

30 Reference is now made to Fig. 3, which is a schematic illustration of a method for analyzing speech of a subject, in accordance with some embodiments of the present invention. The method illustrated in Fig. 3 facilitates the identification of a formant frequency.

Per this method, which is performed in some embodiments, the processor first computes

a cepstrum 48 of each frame, which is the inverse Fourier transform of the logarithm of the magnitude of the Fourier transform of the frame, following phase unwrapping. Cepstrum 48 has both a real and imaginary component. In some embodiments, both of these components are analyzed by the processor, while in other embodiments, the processor analyzes only the real component. The latter is assumed for Fig. 3 and the description below; that is, it is assumed that the "cepstral level" plotted in Fig. 3 is the real component of the cepstrum, and that the processing described below is performed on this real component. (The independent variable of a cepstral graph is called the "quefrency," which is measured in units of samples.)

Subsequently, the processor applies a low-pass cepstral lifter 50 to the cepstrum of the frame. (A filter operating on a cepstrum is referred to as a "cepstral lifter.") A suitable cutoff quefrency for cepstral lifter 50 may be found by requiring that the inverse-transformed liftered cepstrum have a threshold level of smoothness, indicating sufficient liftering-out of the periodic portion of the cepstrum.

Subsequently, the processor inverse-transforms the liftered cepstrum to the frequency domain, such as to derive an inverse-transformed liftered cepstrum signal 52. Then, the processor identifies a formant frequency of the frame, by identifying a peak 54 of signal 52. For example, assuming that peak 54 is located at 350 Hz, the formant frequency of the frame would be 350 Hz. (In the example shown in Fig. 3, the identified formant frequency is the F1 formant. In other cases, the processor may identify other formant frequencies, such as F2, alternatively or additionally to F1.)

As described above, the processor performs the above-described technique for each of the frames of signal 36, and then calculates a single, overall formant frequency from the frame-specific formant frequencies. This overall formant frequency is then used to assess the status of the subject's condition.

Alternatively or additionally, the processor may assess the status of the subject's condition in response to other speech-related parameters extracted from signal 36. Such other parameters may include: a set of mel-frequency cepstrum coefficients (MFCCs) or derivatives thereof, a set of relative spectral transform - perceptual linear prediction (RASTA-PLP) cepstrum components, and/or a set of RASTA-PLP spectrum components. Various other parameters that may be identified from the speech of the subject, and used to assess the subject's condition, include a period of time between breaths, a number of pauses in the speech, loudness of the speech, vocal irregularities, a number of stutters, and an exhalation length, as well as any relevant frequency or pitch parameters. As an example of the latter, the subject's condition may

be assessed based on the subject's inability to produce low frequencies, or based on observed changes in distortion.

Alternatively or additionally, the processor may base the assessment on other, non-speech-related parameters, such as a picture of the subject (which may be acquired, for example, while the subject is speaking, by a camera belonging to the smartphone, or by an external camera), the subject's weight, an echocardiogram (ECG) recording taken from the subject, the subject's arterial oxygen saturation, and/or the subject's level of activity (quantified, for example, as a number of steps taken during a preceding period of time). Any of these parameters may be input to the smartphone, e.g., by the subject. In some embodiments, the smartphone is functionally linked, e.g., wirelessly, to an external device that is configured to measure one or more of these or any other parameters, such as, for example, a camera, a spirometer, a pulse oximeter, an electrocardiograph, a heart rate monitor, a blood pressure monitor, a skin-conductivity sensor, or a posture monitor.

In some embodiments, the processor applies a machine-learned model, such as a classifier or regressor, to some or all of the aforementioned parameters. In such embodiments, the processor is typically also configured to learn the model, based on a plurality of labeled speech recordings. For example, the processor may be provided with a plurality of speech recordings, each labeled as either "wet," indicating that the recording was obtained immediately before dialysis treatment, or "dry," indicating that the recording was obtained immediately after dialysis treatment. Each of these recordings may then be preprocessed, as described above, such as to select the most suitable portion of each recording for parameter (or "feature") extraction. Relevant parameters are then extracted from each recording, and fed to a machine-learning algorithm, such as K Nearest Neighbors (KNN), Decision Trees, or Support Vector Machines (SVM), e.g., with a Gaussian kernel. For example, in some embodiments, the processor extracts, from each frame of each selected portion of each recording, a set of 60 parameters, including (i) a set of 12 MFCCs, (ii) a set of 12 MFCC first derivatives, (iii) a set of 12 MFCC second derivatives, (iv) a set of 8 RASTA-PLP cepstrum components, and (v) a set of 8 RASTA-PLP spectrum components. Typically, a dimensionality reduction algorithm, such as Principal Component Analysis (PCA), is then used reduce the dimensionality of the full set, e.g., to 48 dimensions. This reduced set of parameters is then fed, along with the associated labels, to the relevant machine-learning algorithm. The algorithm then learns a model, which may be applied to subsequent recordings. For example, the algorithm may learn to classify subsequent recordings as either "wet" or "dry."

(In general, assuming D dimensions, F frames per recording, and N recordings, the

"training set" fed to the algorithm includes a matrix of size $D \times FN$. For example, assuming 20 recordings (e.g., 10 "wet" recordings and 10 "dry" recordings), 150 frames per recording, and 48 dimensions, the training set includes a matrix of size 48×3000 .)

In some embodiments, to classify a subsequent recording, the processor first classifies, separately, each frame of the recording. (As described above, the processor may use only the frames belonging to the most suitable portion of the recording.) For example, the processor may apply a learned SVM model, which returns, for each frame, a probability between 0 and 1, with 0 indicating "dry," and 1 indicating "wet." In applying such a model, the processor may set suitable decision thresholds, such that, for example, probabilities below 0.3 deliver a "dry" decision, probabilities above 0.7 deliver a "wet" decision, and intermediate probabilities do not deliver a decision at all. The processor may then compare the total number of "dry" decisions to the total number of "wet" decisions, and apply to the recording the classification that received the greater number of decisions. (In performing this comparison, the processor may give a higher weighting to a probability that is closer to 0 or 1, relative to a probability that is closer to the decision threshold.)

Typically, a separate model is learned for each subject, by acquiring training data from the subject over a suitable period of time, and then applying a suitable machine-learning algorithm to these data.

Although the present disclosure relates mainly to pulmonary conditions associated with the accumulation of excess fluid inside or near the lungs, it is noted that the embodiments described herein may also be applied to other types of pulmonary conditions, such as chronic obstructive pulmonary disease (COPD) or asthma, *mutatis mutandis*.

It will be appreciated by persons skilled in the art that the present invention is not limited to what has been particularly shown and described hereinabove. Rather, the scope of the present invention includes both combinations and subcombinations of the various features described hereinabove, as well as variations and modifications thereof that are not in the prior art, which would occur to persons skilled in the art upon reading the foregoing description.

CLAIMS

1. Apparatus, comprising:
a network interface; and
a processor, configured:
5 to receive, via the network interface, speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid,
 to identify, by analyzing the speech, one or more speech-related parameters of the speech,
 to assess, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and
10 to generate, in response thereto, an output indicative of the status of the pulmonary condition.
2. The apparatus according to claim 1, wherein the condition is selected from the group of conditions consisting of: pulmonary edema, and pleural effusion.
- 15 3. The apparatus according to any one of claims 1-2, wherein the processor is configured to analyze the speech by:
 selecting at least one portion of the speech for analysis, based on one or more metrics selected from the group of metrics consisting of: an energy level of the speech, a stability of a pitch of the speech, a confidence in a calculation of the pitch of the speech, and a stability of a
20 formant frequency of the speech, and
 subsequently, analyzing the selected portion of the speech.
4. The apparatus according to any one of claims 1-3, wherein the processor is configured to assess the status of the pulmonary condition without having identified any wheezing sound in the speech.
- 25 5. A method, comprising:
 using a processor, by analyzing speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid, identifying one or more speech-related parameters of the speech;
 in response to the speech-related parameters, without having identified any cough sounds
30 in the speech, automatically assessing a status of the pulmonary condition; and
 in response to the assessment, generating an output indicative of the status of the pulmonary condition.

6. The method according to claim 5, wherein analyzing the speech comprises performing an analysis selected from the group of analyses consisting of: a spectral analysis of the speech, and a cepstral analysis of the speech.

7. The method according to claim 5, further comprising, by automatically analyzing the speech of the subject, identifying a meaning of the speech, wherein assessing the status comprises assessing the status in response to the meaning.

8. The method according to claim 5, further comprising prompting the subject to provide, by orally responding to a question, the speech.

9. The method according to claim 8, wherein prompting the subject to provide the speech comprises:

automatically placing a call to the subject, and

upon the subject answering the call, using the processor, asking the question to the subject.

10. The method according to claim 5, wherein the condition is selected from the group of conditions consisting of: pulmonary edema, and pleural effusion.

11. The method according to any one of claims 5-10, wherein the speech-related parameters include at least one formant frequency of the speech.

12. The method according to claim 11, wherein analyzing the speech comprises:
for each frame of a time-domain signal that represents the speech:

computing a cepstrum of the frame,

subsequently, applying a low-pass cepstral lifter to the cepstrum of the frame,

subsequently, inverse-transforming the liftered cepstrum to a frequency domain,

and

subsequently, identifying a formant frequency of the frame, by identifying a peak

of the inverse-transformed liftered cepstrum; and

identifying the formant frequency, based on the respective formant frequencies of the frames.

13. The method according to any one of claims 5-10, wherein analyzing the speech comprises:

selecting at least one portion of the speech for analysis, based on one or more metrics selected from the group of metrics consisting of: an energy level of the speech, a stability of a pitch of the speech, a confidence in a calculation of the pitch of the speech, and a stability of a

formant frequency of the speech, and

subsequently, analyzing the selected portion of the speech.

14. The method according to any one of claims 5-10, wherein assessing the status of the pulmonary condition comprises assessing the status of the pulmonary condition by applying a machine-learned model to the speech-related parameters.

15. The method according to any one of claims 5-10, wherein the speech-related parameters include a set of parameters selected from the group of sets consisting of: a set of mel-frequency cepstrum coefficients (MFCCs), a set of derivatives of mel-frequency cepstrum coefficients (MFCCs), a set of relative spectral transform – perceptual linear prediction (RASTA-PLP) cepstrum components, and a set of relative spectral transform – perceptual linear prediction (RASTA-PLP) spectrum components.

16. The method according to any one of claims 5-10, wherein assessing the status of the pulmonary condition comprises assessing the status of the pulmonary condition without identifying any wheezing sound in the speech.

17. The method according to any one of claims 5-7 or claim 10, wherein analyzing the speech of the subject comprises analyzing the speech of a subject without first prompting the subject to provide the speech.

18. A system, comprising:

a mobile communication terminal, comprising a first processor; and

a server, comprising a second processor, the first processor and the second processor being configured to cooperatively carry out a process that includes:

receiving, speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid,

identifying, by analyzing the speech, one or more speech-related parameters of the speech,

assessing, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and

generating, in response to the assessment, an output indicative of the status of the pulmonary condition.

19. A computer software product comprising a tangible non-transitory computer-readable medium in which program instructions are stored, which instructions, when read by a processor, cause the processor:

to receive speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid,

to identify, by analyzing the speech, one or more speech-related parameters of the speech,

to assess, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and

to generate, in response to the assessment, an output indicative of the status of the pulmonary condition.

- 20. Apparatus, comprising:
 - an acoustic sensor; and
 - a processor, configured:

- to receive, via the acoustic sensor, speech of a subject who suffers from a pulmonary condition related to accumulation of excess fluid,

- to identify, by analyzing the speech, one or more speech-related parameters of the speech,

- to assess, in response to the speech-related parameters, a status of the pulmonary condition, without having identified any cough sounds in the speech, and

- to generate, in response to the assessment, an output indicative of the status of the pulmonary condition.




FIG. 1

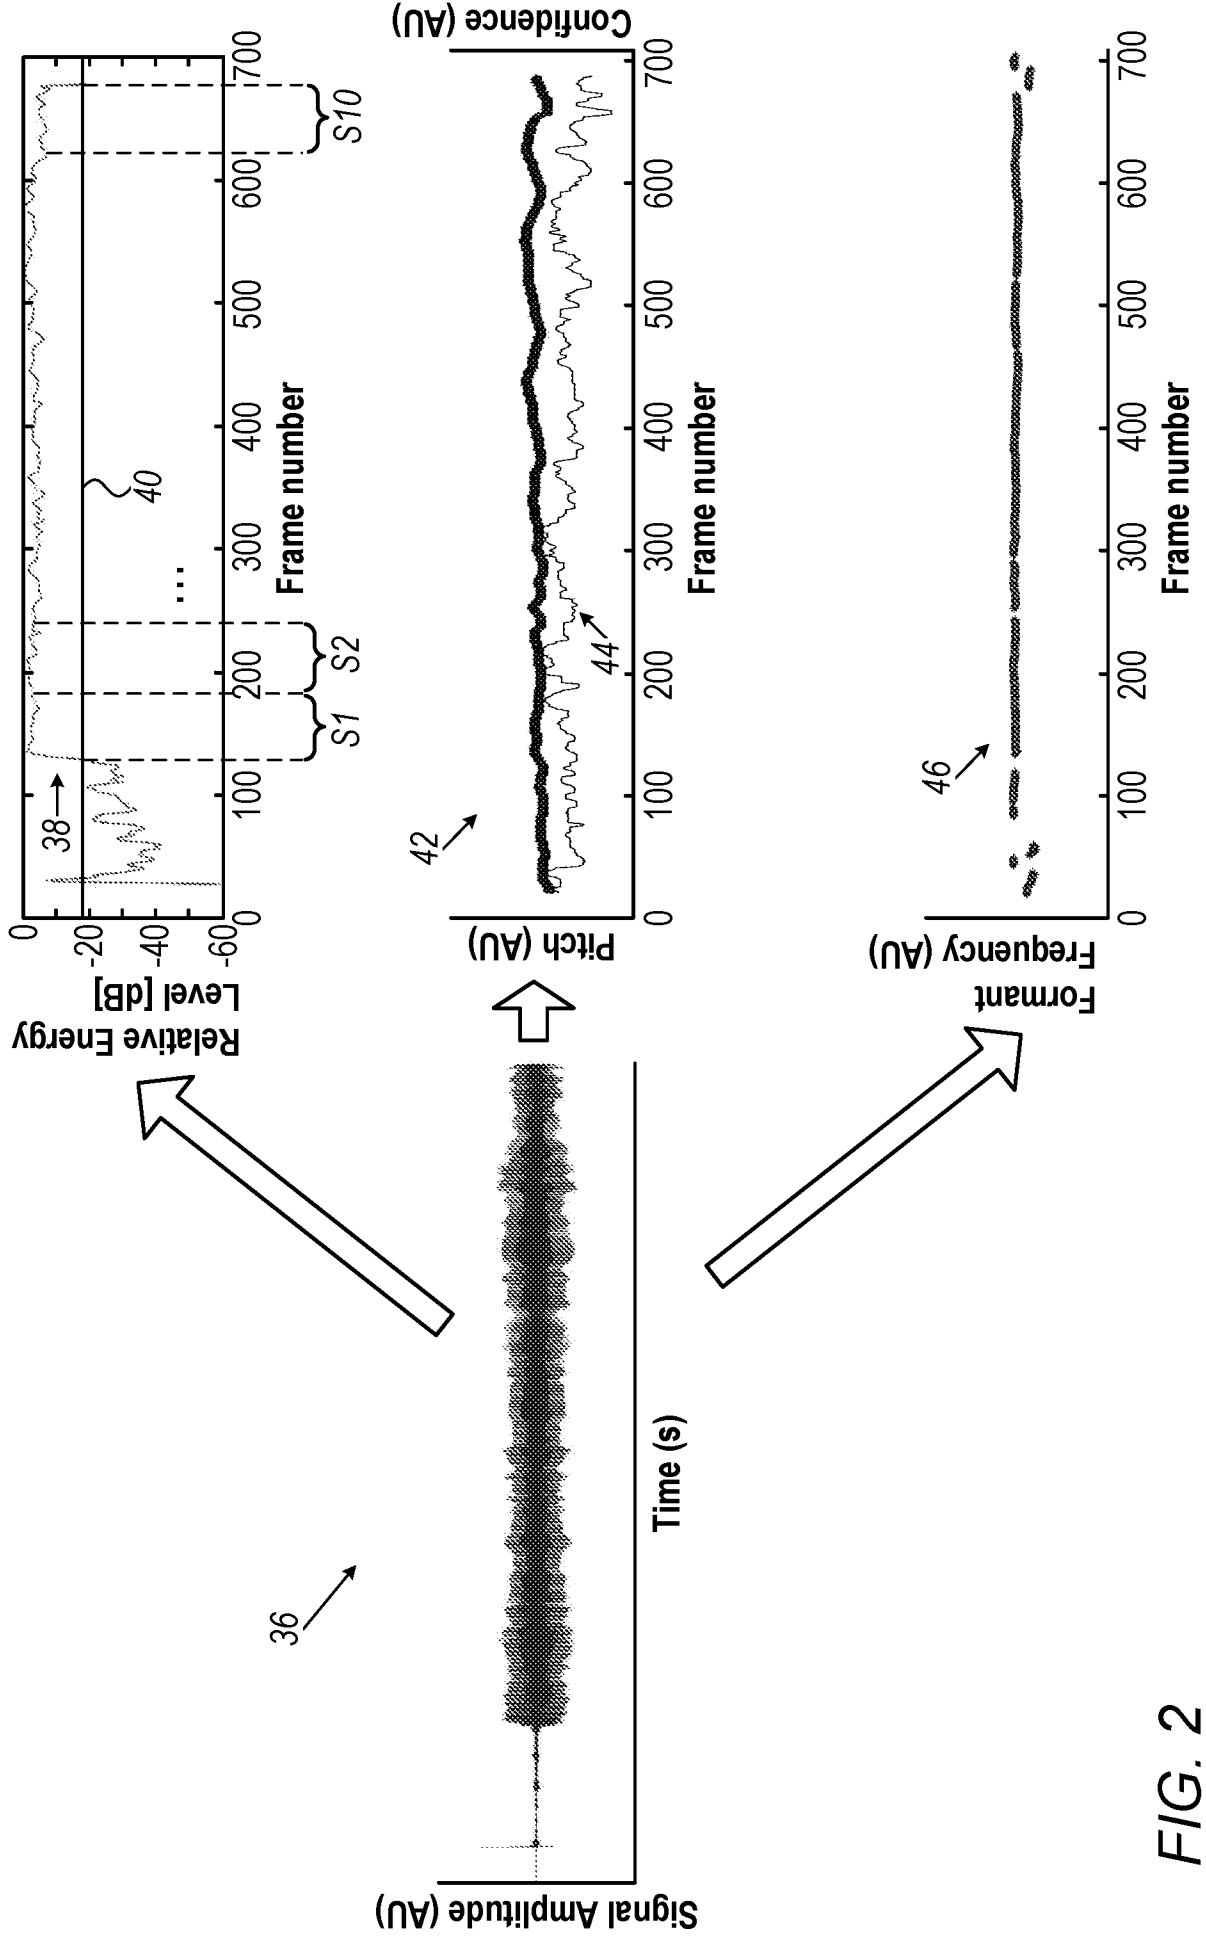


FIG. 2




FIG. 3