



- (51) International Patent Classification:  
*G10L 17/26* (2013.01)      *G10L 25/78* (2013.01)
- (21) International Application Number:  
PCT/US2018/064523
- (22) International Filing Date:  
07 December 2018 (07.12.2018)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
62/595,853      07 December 2017 (07.12.2017) US
- (71) Applicant: **LENA FOUNDATION** [US/US]; 5525 Central Avenue, Boulder, Colorado 80301 (US).
- (72) Inventors: **RICHARDS, Jeffrey**; 5525 Central Avenue, Boulder, Colorado 80301 (US). **HANNON, Stephen**; 5525 Central Avenue, Boulder, Colorado 80301 (US).
- (74) Agent: **SMITH, Cory**; Bryan Cave Leighton Paisner LLP, Two N Central Avenue, Suite 2100, Phoenix, Arizona 85004 (US).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,

(54) Title: SYSTEMS AND METHODS FOR AUTOMATIC DETERMINATION OF INFANT CRY AND DISCRIMINATION OF CRY FROM FUSSINESS

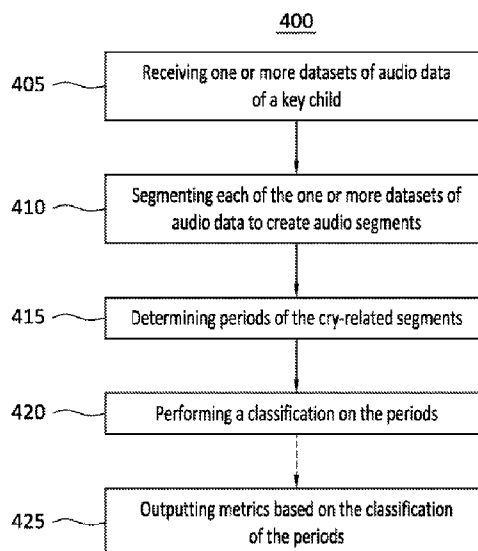


FIG. 4

(57) Abstract: A method including receiving one or more datasets of audio data of a key child captured in a natural sound environment of the key child. The method also includes segmenting each of the one or more datasets of audio data to create audio segments. The audio segments include cry-related segments and non-cry segments. The method additionally includes determining periods of the cry-related segments that satisfy one or more threshold non-sparsity criteria. The method further includes performing a classification on the periods to classify each of the periods as either a cry period or a fussiness period. Other embodiments are described.



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))*

**SYSTEMS AND METHODS FOR AUTOMATIC DETERMINATION OF INFANT CRY  
AND DISCRIMINATION OF CRY FROM FUSSINESS  
CROSS-REFERENCE TO RELATED APPLICATIONS**

**[0001]** This application claims the benefit of U.S. Provisional Application No. 62/595,853, filed December 7, 2017. U.S. Provisional Application No. 62/595,853 is incorporated herein by reference in its entirety.

**TECHNICAL FIELD**

**[0002]** This disclosure relates generally to processing of child digital audio data, and relates more particularly to systems and methods for automatic determination of infant cry and discrimination of cry from fussiness.

**BACKGROUND**

**[0003]** Collection of digitized audio data (e.g., audio datasets from recordings) can be made in a non-invasive fashion in the natural sound environment of a child. These audio datasets can be processed to quantify different types of vocalizations of the child, and to assess the child's language development. These audio datasets also capture instances of cry and/or fussiness from the child. In conventional approaches, assessment of cry or colic in a child has generally involved an imprecise and subjective assessment by parents and/or caregivers.

**BRIEF DESCRIPTION OF THE DRAWINGS**

**[0004]** To facilitate further description of the embodiments, the following drawings are provided in which:

**[0005]** FIG. 1 illustrates a key child in a sound environment wearing an article of clothing that includes a pocket, according to one embodiment;

**[0006]** FIG. 2A illustrates a front view of the pocket of FIG. 1 including an audio capture device, according to an embodiment;

**[0007]** FIG. 2B illustrates a side view of the pocket of FIG. 1 and the audio capture device of FIG. 2A;

**[0008]** FIG. 3 illustrates an audio data processing system including the audio capture device of FIG. 2A and a device for analyzing audio data, according to an embodiment;

**[0009]** FIG. 4 is a flow chart for a method, according to an embodiment;

**[0010]** FIG. 5 illustrates energy levels of sound in a segment associated with the key child of FIG. 1;

**[0011]** FIG. 6 illustrates a cry-related segment plot showing a circular symbol for each cry-related segment of the digital audio dataset based on the elapsed capture time at which the cry-related segment occurs and the cumulative minutes of cry at which the cry-related segment occurs;

[0012] FIG. 7 illustrates a cry-related segment plot in which the cry-related segments from FIG. 6 that are part of periods that satisfy the threshold non-sparsity criteria are labeled with a diamond symbol, and the other cry-related segments FIG. 6 that are filtered out by the non-sparsity criteria are labeled with a circular symbol;

[0013] FIG. 8 illustrates a stacked bar chart showing human-labeled training clips based on the percentage of the duration of the cry (in diagonal), fussiness (in clear), and other (in solid) categories in each of the training clips;

[0014] FIG. 9 illustrates a probability chart showing the predicted value for cry probability across the human-labeled clips, separated into two rows, on an expanded probability axis;

[0015] FIG. 10 illustrates a bubble frequency chart showing the predicted value for cry probability across the human-labeled clips, separated into two rows, on an expanded probability axis, using bubble frequency; and

[0016] FIG. 11 illustrates a cry-related segment plot in which the cry-related segments that are part of periods that satisfy the threshold non-sparsity criteria in FIG. 7 (shown as the diamond symbols in FIG. 7), are classified using the binary logistic regression model as either part of a cry period (and labeled with diamond symbols) or as part of a fussiness period (labeled with square symbols).

[0017] For simplicity and clarity of illustration, the drawing figures illustrate the general manner of construction, and descriptions and details of well-known features and techniques may be omitted to avoid unnecessarily obscuring the present disclosure. Additionally, elements in the drawing figures are not necessarily drawn to scale. For example, the dimensions of some of the elements in the figures may be exaggerated relative to other elements to help improve understanding of embodiments of the present disclosure. The same reference numerals in different figures denote the same elements.

[0018] The terms “first,” “second,” “third,” “fourth,” and the like in the description and in the claims, if any, are used for distinguishing between similar elements and not necessarily for describing a particular sequential or chronological order. It is to be understood that the terms so used are interchangeable under appropriate circumstances such that the embodiments described herein are, for example, capable of operation in sequences other than those illustrated or otherwise described herein. Furthermore, the terms “include,” and “have,” and any variations thereof, are intended to cover a non-exclusive inclusion, such that a process, method, system, article, device, or apparatus that comprises a list of elements is not necessarily limited to those elements but may include other elements not expressly listed or inherent to such process, method, system, article, device, or apparatus.

[0019] The terms “left,” “right,” “front,” “back,” “top,” “bottom,” “over,” “under,” and the like

in the description and in the claims, if any, are used for descriptive purposes and not necessarily for describing permanent relative positions. It is to be understood that the terms so used are interchangeable under appropriate circumstances such that the embodiments of the apparatus, methods, and/or articles of manufacture described herein are, for example, capable of operation in other orientations than those illustrated or otherwise described herein.

**[0020]** As defined herein, “approximately” can, in some embodiments, mean within plus or minus ten percent of the stated value. In other embodiments, “approximately” can mean within plus or minus five percent of the stated value. In further embodiments, “approximately” can mean within plus or minus three percent of the stated value. In yet other embodiments, “approximately” can mean within plus or minus one percent of the stated value.

#### DESCRIPTION OF EXAMPLES OF EMBODIMENTS

**[0021]** A number of embodiments include a system including one or more processors and one or more non-transitory computer-readable media storing computing instructions configured to run on the one or more processors and perform various acts. The acts can include receiving one or more datasets of audio data of a key child captured in a natural sound environment of the key child. The acts also can include segmenting each of the one or more datasets of audio data to create audio segments. The audio segments can include cry-related segments and non-cry segments. The acts additionally can include determining periods of the cry-related segments that satisfy one or more threshold non-sparsity criteria. The acts further can include performing a classification on the periods to classify each of the periods as either a cry period or a fussiness period.

**[0022]** Various embodiments include a method being implemented via execution of computing instructions configured to run on one or more processors and stored on one or more non-transitory computer-readable media. The method can include receiving one or more datasets of audio data of a key child captured in a natural sound environment of the key child. The method also can include segmenting each of the one or more datasets of audio data to create audio segments. The audio segments can include cry-related segments and non-cry segments. The method additionally can include determining periods of the cry-related segments that satisfy one or more threshold non-sparsity criteria. The method further can include performing a classification on the periods to classify each of the periods as either a cry period or a fussiness period.

**[0023]** Infantile colic (or “colic”) generally involves significant amounts of uncontrolled and/or inconsolable crying in an infant. In some definitions, colic is defined as episodes of crying for more than three hours a day, for more than three days a week, for more than three weeks for an otherwise healthy infant who is not ill, hungry, uncomfortable, in obvious pain, etc. The infant

is typically less than four months of age. The crying involved in colic is generally high-intensity crying, as opposed to fussiness, which involves lower-intensity crying. For example, in fussiness, the infant typically has regular breathing with normal air intake between cries, but in colicky crying, the infant's breathing often involves agitated breathing with sudden air intake between cries. As another example, high-intensity cries are often louder and higher pitched than cries in fussiness.

**[0024]** In conventional approaches, determining the amount of crying from an infant, and/or determining whether the infant has colic, is typically done through parental reporting. Parental reporting, however, has several disadvantages. For example, much of parental reporting is based on memory, which is unreliable. In many cases, parents are instructed to use a baby diary log to record when the infant cries or is inconsolable, but most parents have a hard time regularly recording all episodes of crying with accuracy. One reason is that a crying child typically generates a great deal of stress for parents. Further, the parents often have difficulty distinguishing between intense crying and non-intense crying. Even when parental recording is performed regularly, there are subjective differences between the reporting of different parents.

**[0025]** In order to determine if an intervention (e.g., medicine, a change in diet, etc.) is effective, it is advantageous to be able to quantify the amount of cry, as opposed to fussiness, from the infant. By quantifying the amount of cry in an objective way, the subjective differences between parents in parental reporting can be overcome, beneficially providing an objective consistency in quantifying the level of crying and/or colic.

**[0026]** In a number of embodiments, an audio capture device can be worn by a key child to capture datasets of audio data of the natural sound environment of the child in a non-invasive fashion. These datasets of audio data can be processed into segments to quantify aspects of the datasets, such as to monitor and/or analyze the natural sound environment, vocalizations, and the development of a key child. A key child as used herein may be a child, an adult, such as an adult with developmental disabilities, or any individual who produces sounds that are of interest. A natural sound environment of the key child can be monitored without placing artificial limitations on the activities of the key child or requiring a third-party observer. The natural sound environment can be analyzed to identify the types of sounds produced by the key child.

**[0027]** A natural sound environment can include an area surrounding a key child during his or her normal daily activities and contain sources of sounds that may include the key child, other children, an adult, an electronic device, and background noise. A clinical or research environment, by contrast, can include a controlled environment or location that contains pre-selected or natural sources of sounds.

**[0028]** In some embodiments, the key child may wear an article of clothing that includes an

audio capture device located in a pocket attached to or integrated with the article of clothing. The audio capture device may be configured to record and store, and/or otherwise process, audio data associated with the sound environment of the key child for a predetermined amount of time. The datasets of audio data can include background or environmental noise, silence, the spoken words or other sounds of the key child, words spoken by others, sounds from electronic devices such as televisions and radios, or any sound or words from any source. The location of the audio capture device preferably allows it to capture the vocal productions of the key child without interfering in the normal activities of the key child. During or after the pre-set amount of time, the datasets of audio data captured by the audio capture device can be processed and analyzed. For example, the datasets of audio data may be analyzed to identify segments and assign a segment ID or a source for each audio segment using a Minimum Duration Gaussian Mixture Model (MD-GMM). In some embodiments, the processing can be performed in a processing device outside the audio data. In other embodiments, the audio data captured can be processed in the audio capture device. For example, in various embodiments, the audio data can be processed as it is received.

**[0029]** Sources for each audio segment can include the key child, an adult, another child, an electronic device, or any person or object capable of producing sounds. Sources may also include general sources that are not associated with a particular person or device. Examples of such general sources include background or environmental noise, silence, and overlapping sounds. In some embodiments, sources are identified by analyzing each audio segment using models of different types of sources. The models may include audio characteristics commonly associated with each source. Audio segments for which the key child or an adult is identified as the source may be further analyzed, such as by determining certain characteristics associated with the key child and/or adult, to provide metrics associated with the sounds produced by the key child.

**[0030]** In some embodiments, an audio capture device may be used to capture, process, record, and/or store audio associated with the sound environment of the key child. The audio capture device may be any type of device adapted to process audio data and/or capture and store audio data. The audio capture device can store audio over a long term (e.g., a day) or a short term (e.g., brief storage, such as a few seconds, in a memory buffer). The audio capture device can be located in or around a sound environment of a child. In some embodiments, the audio capture device can include one or more microphones connected to a storage device and can be located in one or more rooms that the key child often occupies. In other embodiments, the audio capture device is located in an article of clothing worn by the child.

**[0031]** Turning to the drawings, FIG. 1 illustrates a key child, such as child 100, in a sound

environment 102 wearing an article of clothing 104 that includes a pocket 106. The pocket 106 may include an audio capture device (not shown) that is adapted to capture audio from sound environment 102. Sound environment 102 can be an area surrounding child 100 that includes sources for audio (not shown), including one or more adults, other children, and/or electronic devices such as a television, a radio, a toy, background noise, or any other source that produces sounds. Examples of sound environment 102 include a natural sound environment or a clinical or research environment. Article of clothing 104 can be the normal clothing of child 100, a vest over the normal clothing of child 100, or any article of clothing commonly worn by child 100.

**[0032]** In some embodiments, the audio capture device is placed at or near the center of the chest of the key child. However, other placements are possible. The audio capture device in pocket 106 can be any device capable of capturing audio associated with the sound environment of child 100.

**[0033]** One example of an audio capture device is a digital audio recorder of the LENA system of the LENA Foundation in Boulder, Colorado. The audio capture device can be relatively small and lightweight and can be placed in pocket 106. Pocket 106 can hold the audio capture device in place in an unobtrusive manner so that the audio capture device does not distract child 100, other children, and adults that interact with child 100.

**[0034]** Turning ahead in the drawings, FIG. 2A illustrates a front view of pocket 106 including an audio capture device 108, and FIG. 2B illustrates a side view of pocket 106 including audio capture device 108. Audio capture device 108 is merely exemplary, and embodiments of the system are not limited to the embodiments presented herein. The audio capture device can be employed in many different embodiments or examples not specifically depicted or described herein. Pocket 106 can be designed to keep audio capture device 108 in place and to minimize acoustic interference. Pocket 106 can include an inner area 110 formed by a main body 112 and an overlay 114 connected to main body 112 via stitches 116 or another suitable connecting mechanism. Main body 112 can be part of article of clothing 104, or alternatively attached to article of clothing 104 (FIG. 1) using stitches or by another suitable connecting mechanism. A stretch layer 118 can be located in inner area 110 and attached to main body 112 and overlay 114 via stitches 116. Audio capture device 108 can be located between main body 112 and stretch layer 118. Stretch layer 118 can be made of a fabric adapted to stretch but provide a force against audio capture device 108 to retain audio capture device 108 in its position. For example, stretch layer 118 can be made from a blend of nylon and spandex, such as eighty-five percent (85%) nylon and fifteen percent (15%) spandex, which can help to keep audio capture device 108 in place. Overlay 114 can cover stretch layer 118 and, in some embodiments, can include at least one opening in which the microphone of audio capture device 108 is located.



The opening can be covered with a material that provides certain desired acoustic properties. In one embodiment, the material is one hundred percent (100%) cotton.

**[0035]** Pocket 106 also can include snap connectors 120 by which overlay 114 is opened and closed to install or remove the audio capture device 108. In some embodiments, at least one of stitches 116 can be replaced with a zipper to provide access to the audio capture device 108 in addition to or as an alternative to using snap connectors 120.

**[0036]** When audio capture device 108 includes multiple microphones, pocket 106 can include multiple openings that correspond to the placement of the microphones on audio capture device 108. The particular dimensions of pocket 106 can change as the design of audio capture device 108 changes or as the number or type of microphones change. In some embodiments, pocket 106 can position the microphone relative to the mouth of the key child to provide certain acoustical properties and secure the microphone (and optionally audio capture device 108) in a manner that does not result in friction noises captured by the microphone. Audio capture device 108 can be turned on and thereafter capture audio, including sounds produced by the key child, other children, and adults, as well as other types of sounds that the child encounters, including television, toys, environmental noises, etc. The audio data can be stored in audio capture device 108. In some embodiments, audio capture device 108 can be periodically removed from pocket 106 and the stored audio can be analyzed, or the audio can be transmitted wirelessly from audio capture device 108. In other embodiments, the data can be processed within audio capture device 108, such as while the data is being captured.

**[0037]** Turning ahead in the drawings, FIG. 3 illustrates an audio data processing system including audio capture device 108 and a device 200 for analyzing datasets of audio data. Device 200 is merely exemplary and embodiments of the system are not limited to the embodiments presented herein. The device can be employed in many different embodiments or examples not specifically depicted or described herein. In some embodiments, certain elements or modules of device 200 can perform various procedures, processes, and/or activities. In other embodiments, the procedures, processes, and/or activities can be performed by other suitable elements or modules of device 200.

**[0038]** In many embodiments, device 200 can include a processor 202 and a computer-readable medium, such as memory 204. Audio capture device 108 may be connected to device 200 via wireline or wirelessly. In some embodiments, audio capture device 108 can be connected to device 200 via a USB cable. Device 200 can be any type of processor-based device, examples of which include a computer and a server. For example, in some embodiments, device 200 can be a general-purpose mobile device, such as a smartphone. In other embodiments, device 200 can be specific hardware configured for implementing at least a portion of the systems and/or

methods described herein. Memory 204 can be configured to store computer-executable instruction code and data. The computer-executable instruction code can include an application 206, such as a data analysis application that can be used to view, generate, and output data analysis. Application 206 can include an audio engine 208 that, as described in more detail below, can be configured to perform methods according to various embodiments to analyze audio data and generate metrics associated therewith. In some embodiments, audio engine 208 can be a separate application that is executable separate from, and optionally concurrent with, application 206. Memory 204 also can include a data storage 210 that is configured to store data generated by the application 206 or audio engine 208, or input by a user. In some embodiments, data storage 210 can be separate from device 200 but connected to device 200 via wireline or wireless connection.

**[0039]** In a number of embodiments, device 200 can be in communication with an input device 212 and an output device 214. Input device 212 can be configured to receive user input and communicate the user input to device 200. Examples of input device 212 can include a keyboard, mouse, scanner, and network connection. User inputs can include commands that cause processor 202 to execute various functions associated with application 206 or audio engine 208. In many embodiments, output device 214 can be configured to provide data or visual output from application 206 or audio engine 208. In some embodiments, output device 214 can display a graphical user interface (GUI) that includes one or more selectable buttons that are associated with various functions provided by application 206 or audio engine 208. Examples of output device 214 can include a monitor, network connection, and printer. Input device 212 can be used to setup or otherwise configure audio engine 208. For example, the age of the key child and other information associated with the key child's sound environment may be provided to audio engine 208 and stored in local storage 210 during a setup or configuration via input device 212.

**[0040]** In many embodiments, the audio file stored on audio capture device 108 can be uploaded to device 200 and stored in local storage 210. In some embodiments, the audio data can be uploaded in a proprietary format which can prevent the playback of the speech from device 200 or access to content of the speech, thereby promoting identity protection of the speakers. In other embodiments, the audio data can be uploaded without being encoded to allow for storage in local storage 210 and playback of the file or portions of the file.

**[0041]** In some embodiments, device 200 can be a web server, and input device 212 and output device 214 can be combined to form a computer system that sends data to and receives data from device 200 via a network connection. Input device 212 and output device 214 can be used to access application 206 and audio engine 208 remotely and cause it to perform various

functions according to various embodiments. Audio capture device 108 can be connected to input device 212 and output device 214, and the audio data stored on audio capture device 108 can be uploaded to device 200 over a network such as the Internet or an intranet in which the audio data are processed and metrics are provided to output device 214. In some embodiments, the audio data received remotely from input device 212 and output device 214 can be stored in local storage 210 and subsequently accessed for research purposes or otherwise.

**[0042]** To reduce the amount of memory needed on audio capture device 108, the audio data can be compressed. In some embodiments, a digital visual interface (DVI), such as a DVI-4 adaptive differential pulse-code modulation (ADPCM) compression scheme can be used. If a compression scheme is used, then the file can be decompressed after it is uploaded to device 200 to a normal linear pulse code modulation (PCM) audio format or other suitable audio format.

**[0043]** Turning ahead in the drawings, FIG. 4 illustrates a flow chart for a method 400, according to an embodiment. In many embodiments, method 400 can be a method of determining cry segments and discrimination of cry from fussiness. Method 400 is merely exemplary and is not limited to the embodiments presented herein. Method 400 can be employed in many different embodiments or examples not specifically depicted or described herein. In some embodiments, the procedures, the processes, and/or the activities of method 400 can be performed in the order presented. In other embodiments, the procedures, the processes, and/or the activities of method 400 can be performed in any suitable order. In still other embodiments, one or more of the procedures, the processes, and/or the activities of method 400 can be combined or skipped. In some embodiments, method 400 can be performed by device 200 (FIG. 3).

**[0044]** Referring to FIG. 4, method 400 can include a block 405 of receiving one or more datasets of audio data of a key child captured in a natural sound environment of the key child. In many embodiments, the datasets of audio data can be similar or identical to the audio data captured by audio capture device 108 (FIGs. 2A, 2B, 3). In some embodiments, the datasets of audio data can be digitized audio data, such as digital audio recordings. In many embodiments, the datasets of audio data can be received by being uploaded into device 200 (FIG. 3), as described above. For example, the datasets of audio data can be uploaded to a cloud-based server. In other embodiments, the datasets of audio data can be captured and processed on audio capture device 108 (FIGs. 2A, 2B, 3) without uploading to another device. In a number of embodiments, the datasets of audio data of the key child can be datasets of audio data in which the key child is less than 4 months old. For example, the datasets of audio data can be used to quantify cry in infantile colic. In other embodiments, the datasets of audio data can be datasets of audio data in which the key child is another suitable age. For example, in certain other

embodiments, the datasets of audio data can be used to quantify cry for other purposes, such as part of analyzing a child for autism spectrum disorder (ASD), and the key child can be, for example, between 1 to 3 years old. The datasets of audio data can capture the natural sound environment of the key child over a suitable amount of time, such as 8 hours, 12 hours, 16 hours, 24 hours, or another suitable duration.

**[0045]** In a number of embodiments, method 400 also can include a block 410 of segmenting each of the one or more datasets of audio data to create a plurality of audio segments. In many embodiments, the plurality of audio segments can include a plurality of cry-related segments and a plurality of non-cry segments. In many embodiments, the segmenting of the datasets of audio data can create segments having a duration of approximately 300 milliseconds to 1.0 second in length. In some embodiments, the duration of the segments can be variable across the segments. In other embodiments, the duration of the segments can be fixed across all of the segments. As an example, the segmenting of block 410 can be performed by device 200 (FIG. 3).

**[0046]** In many embodiments, the segmenting process can involve using machine-learning algorithms to automatically parse the datasets of audio data into labeled time segments. For example, each of the labeled segments can fall into one of eight categories that indicates the dominant source of sound in the segment, namely: (1) male adult, (2) female adult, (3) key child, (4) other child, (5) overlap, (6) background or environmental noise, (7) TV/electronic, and (8) silence. The key child is the child wearing the audio capture device (e.g., 108 (FIGs. 2A, 2B, 3)). Although there can be multiple co-occurring sounds within a time segment, such as a mother talking while the key child is crying and a TV is on in the background, in some embodiments, the segment processing can label the segment based on the sound that is most dominant or has the highest likelihood of the eight categories. For example, the mother may be attempting to soothe the key child with a song, and when the baby has an intake of breath, the mother's voice can be dominant over the cries of the key child. For the sounds labeled as key child, the segmenting process can involve assigning subcategories, namely: (1) VOC for vocalization sounds, (2) VEG for vegetative sounds, and (3) CRY for cry-related sounds.

**[0047]** In some embodiments, to perform the segmenting process, an audio engine (e.g., audio engine 208 (FIG. 3)) can divide the dataset of audio data into one or more audio segments and can identify a segment ID or source for each of the audio segments from the dataset of audio data received from the audio capture device 108 (FIGs. 2A, 2B, 3). This process is referred to herein as "segmentation" and "segment ID." In some embodiments, an audio segment can be a portion of the dataset of audio data (e.g., a portion of an audio recording) having a certain duration and including acoustic features associated with the natural sound environment of the key child during the duration. The dataset of audio data can include a number of audio

segments, each associated with a segment ID or source. Sources can include an individual or a device that produces the sounds within the audio segment. For example, an audio segment may include the sounds produced by the key child, who is identified as the source for that audio segment. Sources also can include other children, adults, electronic devices, background or environmental noise, overlapped sounds, and silence. Electronic devices can include televisions, radios, telephones, toys, and any device that provides recorded or simulated sounds such as human speech.

**[0048]** Sources associated with each of the audio segments can be identified to assist in further classifying and analyzing the dataset of audio data. Some metrics provided by some embodiments include data regarding certain sources and disregard data from other sources. For example, audio segments associated with live speech directed to the key child can be distinguished from audio segments associated with electronic devices.

**[0049]** To perform segmentation to generate the audio segments and identify the sources for each segment, a number of models may be used that correspond to the key child, other children, male adult, female adult, background noise, TV sound, silence, and overlap. Alternative embodiments may use more, fewer, or different models to perform segmentation and identify a corresponding segment ID. One such technique performs segmentation and segment ID separately. Another technique performs segmentation and identifies a segment ID for each segment concurrently.

**[0050]** Traditionally, a Hidden Markov Model (HMM) with minimum duration constraint has been used to perform segmentation and identify segment IDs concurrently. A number of HMM models may be provided, each corresponding to one source. The result of the model can be a sequence of sources with a likelihood score associated with each source. The optimal sequence can be searched using a Viterbi algorithm or dynamic programming and the “best” source can be identified for each segment based on the score. However, this approach can be complex for some segments in part because it uses transition probabilities from one segment to another, that is, the transition between each segment. Transition probabilities are related to duration modeling of each source. HMM duration models can have discrete geometric distribution or continuous exponential distribution, which may not be appropriate for the sound sources of concern. Most audio data can include segments having a high degree of variation in their duration. Although the HMM model can be used in some embodiments, alternative techniques can be used to perform segmentation and segment ID.

**[0051]** An alternative technique used in some embodiments to perform segmentation and segment ID is a Minimum Duration Gaussian Mixture Model (MD-GMM). Each model of the MD-GMM can include criteria or characteristics associated with sounds from different sources.

Examples of models of the MD-GMM include a key child model that includes characteristics of sounds from a key child, an adult model that includes characteristics of sounds from an adult, an electronic device model that includes characteristics of sounds from an electronic device, a noise model that includes characteristics of sounds attributable to background or environmental noise, another child model that includes characteristics of sounds from a child other than the key child, a parentese model that includes complexity level speech criteria of adult sounds, an age-dependent key child model that includes characteristics of sounds of a key child of different ages, and a loudness/clearness detection model that includes characteristics of sounds directed to a key child. Some models can include additional models. For example, the adult model can include an adult male model that includes characteristics of sounds of an adult male and an adult female model that includes characteristics of sounds of an adult female. The models can be used to determine the source of sound in each segment by comparing the sound in each segment to criteria of each model and determining if a match of a pre-set accuracy exists for one or more of the models.

**[0052]** In some embodiments, the MD-GMM technique begins when the dataset of audio data (e.g., an audio recording) is converted to a sequence of frames or segments. Segments having a duration of  $2*D$ , where  $D$  is a minimum duration constraint, are identified using a maximum log-likelihood algorithm for each type of source. The maximum score for each segment is identified. The source associated with the maximum score is correlated to the segment for each identified segment.

**[0053]** In many embodiments, audio engine 208 (FIG. 3) can process datasets of audio data using the maximum likelihood MD-GMM to perform segmentation and segment ID. Audio engine 208 (FIG. 3) can search all possible segment sequences under a minimum duration constraint to identify the segment sequence with maximum likelihood. One possible advantage of MD-GMM is that any segment longer than twice the minimum duration ( $2*D$ ) could be equivalently broken down into several segments with a duration between the minimum duration ( $D$ ) and two times the minimum duration ( $2*D$ ), such that the maximum likelihood search process ignores all segments longer than  $2*D$ . This can reduce the search space and processing time. One implementation of using maximum likelihood MD-GMM is explained below. Other implementations are also possible, such as:

1. Acoustic Feature Extraction – the audio stream is converted to a stream of feature vectors  $\{X_1, X_2, \dots, X_T \mid X_i \in R^n\}$  using a feature extraction algorithm, such as the MFCC (mel-frequency cepstrum coefficients).

2. Log likelihood calculation for a segment  $\{X_1, X_2, \dots, X_S\}$ :

$$Lcs = \sum_{i=1}^S \log(f_c(X_i)),$$

where  $f_c(X_i)$  is the likelihood of frame  $X_i$  being in class  $c$ .

**[0054]** The following describes one procedure of maximum likelihood MD-GMM search:

1. Initialize searching variables:  $S(c,0,0) = 0$ ,  $c = 1, \dots, C$ , where  $c$  is the index for all segment classes. Generally, the searching variable  $S(c,b,n)$  represents the maximum log-likelihood for the segment sequence up to the frame  $b-1$  plus the log-likelihood of the segment from frame  $b$  to frame  $n$  being in class  $c$ .

2. Score frames for  $n = 1, \dots, T$ , i.e. all feature frames:

$$S(c,b,n) = S(c,b,n-1) + \log(f_c(X_n), \quad \forall b,c,n-b < 2^*D_c,$$

i.e., the current score at frame  $n$  could be derived from the previous score at frame  $n-1$ .

The searching variable for segments less than twice the minimum duration is retained.

3. Retain a record of the optimal result at frame  $n$  (similarly, segments under twice the minimum duration will be considered):

$$S^*(n) = \max_{c,b,2^*D_c > n-b > D_c} S(c,b,n)$$

$$B^*(n) = \arg \max_{b,(c,b,2^*D_c > (n-b) > D_c} S(c,b,n)$$

$$C^*(n) = \arg \max_{c,(c,b,2^*D_c > (n-b) > D_c} S(c,b,n)$$

4. Initialize new searching variables for segments starting at frame  $n$ :

$$S(c,n,n) = S^*(n), \quad \forall c$$

5. Iterate step 4 to step 6 until the last frame  $T$ .

6. Trace back to get the maximum likelihood segment sequence.

The very last segment of the maximum likelihood segment sequence is  $(C^*(T), B^*(T), T)$ , that is, the segment starting from frame  $B^*(T)$  and ending with frame  $T$  with class id of  $C^*(T)$ . The remaining segments can be obtained in the best sequence by using the following back-tracing procedure:

6.1. Initialize back-tracing:

$$t = T, m = 1$$

$$S(m) = (C^*(t), B^*(t), t)$$

6.2. Iterate back-tracing until  $t=0$

$$C\_current = C^*(t)$$

$$t = B^*(t)$$

If  $C^*(t) = C\_current$ , then do nothing;

Otherwise,  $m = m + 1$ ,  $S(m) = (C^*(t), B^*(t), t)$

**[0055]** Additional processing can be performed to further refine identification of segments associated with the key child or an adult as sources. As stated above, the language environment can include a variety of sources that may be identified initially as the key child or an adult when the source is actually a different person or device. For example, sounds from a child other than the key child may be initially identified as sounds from the key child. Sounds from an electronic device may be confused with live speech from an adult. Furthermore, some adult sounds may be detected that are directed to another person other than the key child. Certain embodiments may implement methods for further processing and refining the segmentation and segment ID to decrease or eliminate inaccurate source identifications and to identify adult speech directed to the key child.

**[0056]** Further processing can occur concurrently with, or subsequent to, the initial MD-GMM model described above. One embodiment is an adaptation method for further processing the dataset of audio data by modifying models associated with the MD-GMM subsequent to an initial MD-GMM. First, audio engine 208 (FIG. 3) can process the dataset of audio data using a first MD-GMM. For example, the dataset of audio data is processed in accordance with the MD-GMM described above to perform an initial segmentation and segment ID.

**[0057]** Second, audio engine 208 (FIG. 3) can modify at least one model of the MD-GMM. Audio engine 208 (FIG. 3) can automatically select one or more models of the MD-GMM to modify based on pre-set steps. In some embodiments, if audio engine 208 (FIG. 3) detects certain types of segments that involve further scrutiny, it can select the model of the MD-GMM that is most related to the types of segments detected to modify (or for modification). Any model associated with the MD-GMM can be modified. Examples of models that can be modified include the key child model with an age-dependent key child model, an electronic device model, a loudness/clarity model that may further modify the key child model and/or the adult model, and a parentese model that may further modify the key child model and/or the adult model.

**[0058]** Third, audio engine 208 (FIG. 3) can process the dataset of audio data again using the modified models of the MD-GMM. The second process can result in a different segmentation and/or segment ID based on the modified models, providing a more accurate identification of the source associated with each segment.

**[0059]** Fourth, audio engine 208 (FIG. 3) can determine whether to perform additional model



modification. In some embodiments, audio engine 208 (FIG. 3) analyzes the new segmentation and/or segment ID to determine if any segments or groups of segments would benefit from additional scrutiny. In some embodiments, audio engine 208 (FIG. 3) can access data associated with the language environment in data storage 210 (FIG. 3) and can use it to determine if additional model modification would be beneficial, such as a modification of the key child model based on the current age of the child. If additional model modification would be beneficial, the process can return to step 2 of modifying the model of the MD-GMM described above for additional MD-GMM model modification. If no additional model modification would be beneficial, the process can proceed to analyzing the segment sound. Certain embodiments of modifying exemplary models in accordance with various embodiments are described below. Other models than those described below may be modified in certain embodiments.

**[0060]** In some embodiments, audio engine 208 (FIG. 3) can implement an age-dependent key child model concurrently with, or subsequent to, the initial MD-GMM to modify the key child model of the MD-GMM to more accurately distinguish segments in which other children are the source from segments in which the key child is the source. For example, the MD-GMM can be modified to implement an age-dependent key child model during the initial or a subsequent segmentation and segment ID.

**[0061]** The key child model can be age dependent because the audio characteristics of the vocalizations, including utterances and other sounds, of a key child can change dramatically over the time that audio capture device 108 (FIGs. 2A, 2B, 3) can be used. Although the use of two separate models within the MD-GMM, one for the key child and one for other children, can identify the vocalizations of the key child, the use of an age-dependent key child model further helps to reduce the confusion between vocalizations of the key child and vocalizations of the other children. In one embodiment, the age-dependent key child models are: (1) less than one-year old, (2) one-year old, (3) two-years old, and (4) three-years old. Alternative embodiments can use other age groupings and/or may use groupings of different age groups. For example, other embodiments could use monthly age groups or a combination of monthly and yearly age groups. Each of the models can include characteristics associated with sounds commonly identified with children of the age group.

**[0062]** In one embodiment, the age of the key child can be provided to device 200 (FIG. 3) via input device 212 (FIG. 3) during a set-up or configuration. Audio engine 208 (FIG. 3) can receive the age of the key child and can select one or more of the key child models based on the age of the key child. For example, if the key child is one year and ten months old, audio engine 208 can select key child model 2 (one-year-old model) and key child model 3 (two-years-old model) or only key child model 2 based on the age of the key child. Audio engine 208 (FIG. 3)

can implement the selected key child model or models by modifying the MD-GMM models to perform the initial or a subsequent segmentation and segment ID.

**[0063]** In some embodiments, any segments including sounds, such as words or speech, that are generated electronically by an electronic device can be identified as such, as opposed to an inaccurate identification as live speech produced by an adult. Electronic devices can include a television, radio, telephone, audio system, toy, or any electronic device that produces recordings or simulated human speech. In some embodiments, audio engine 208 (FIG. 3) can modify an electronic device model in the MD-GMM to more accurately identify segments from an electronic device source and separate them from segments from a live adult without determining the content of the segments and without limiting the environment of the speaker (e.g., without requiring the removal of or inactivation of the electronic devices from the language environment).

**[0064]** In many embodiments, audio engine 208 (FIG. 3) can be adapted to modify and use the modified electronic device model concurrently with, or subsequent to, the initial MD-GMM process. In some embodiments, the electronic device model can be implemented after a first MD-GMM process is performed and used to adapt the MD-GMM for additional determinations using the MD-GMM for the same dataset of audio data. In a number of embodiments, audio engine 208 (FIG. 3) can examine segments segmented using a first MD-GMM to further identify reliable electronic segments. Reliable electronic segments can be segments that are more likely associated with a source that is an electronic device and include certain criteria. For example, audio engine 208 (FIG. 3) can determine if one or more segments include criteria commonly associated with sounds from electronic devices. In some embodiments, the criteria can include (1) a segment that is longer than a predetermined period or is louder than a predetermined threshold; or (2) a series of segments having a pre-set source pattern. An example of one predetermined period is five seconds. An example of one pre-set source pattern can include the following:

Segment 1 – Electronic device source;

Segment 2 – A source other than the electronic device source (e.g., adult);

Segment 3 – Electronic device source;

Segment 4 – A source other than the electronic device source; and

Segment 5 – Electronic device source.

**[0065]** The reliable electronic device segments can be used to adapt the MD-GMM to include an adaptive electronic device model for further processing. For example, audio engine 208 (FIG. 3)

can use a regular K-means algorithm as an initial model and tune it with an expectation-maximization (EM) algorithm. The number of Gaussians in the adaptive electronic device model can be proportional to the amount of feedback electronic device data and not exceed an upper limit. In one embodiment, the upper limit is 128. Other suitable upper limits can be used.

**[0066]** In several embodiments, audio engine 208 (FIG. 3) can perform the MD-GMM again by applying the adaptive electronic device model to each frame of the sequence to determine a new adaptive electronic device log-likelihood score for frames associated with a source that is an electronic device. The new score may be compared with previously stored log-likelihood scores for those frames. Audio engine 208 (FIG. 3) can select the larger log-likelihood score based on the comparison. The larger log-likelihood score can be used to determine the segment ID for those frames.

**[0067]** In some embodiments, the MD-GMM modification using the adaptive electronic device model can be applied using a pre-set number of consecutive equal length adaptation windows moving over all frames. The audio data can be divided into overlapping frames having a pre-set length. An example of frame length according to one embodiment is 25.6 milliseconds with a 10 millisecond shift resulting in 15.6 milliseconds of frame overlap. The adaptive electronic device model can use local data obtained using the pre-set number of adaptation windows. An adaptation window size of 30 minutes can be used in some embodiments. An example of one pre-set number of consecutive equal length adaptation windows is three. In some embodiments, adaptation window movement does not overlap. The frames within each adaptation window can be analyzed to extract a vector of features for later use in statistical analysis, modeling, and classification algorithms. The adaptive electronic device model can be repeated to further modify the MD-GMM process. For example, the process may be repeated three times.

**[0068]** In order to select the frames that are most useful for identifying the speaker, some embodiments use frame level near/far detection or loudness/clearness detection model. Loudness/clearness detection models can be performed using a Likelihood Ratio Test (LRT) after an initial MD-GMM process is performed. At the frame level, the LRT can be used to identify and discard frames that could confuse the identification process. For each frame, the likelihood for each model can be calculated. The difference between the most probable model likelihood and the likelihood for the silence model is calculated and the difference is compared to a predetermined threshold. Based on the comparison, the frame is either dropped or used for segment ID. For example, if the difference meets or exceeds the predetermined threshold, then the frame is used; but if the difference is less than the predetermined threshold, then the frame is dropped. In some embodiments, frames are weighted according to the LRT.

**[0069]** In many embodiments, audio engine 208 (FIG. 3) can use the LRT to identify segments

directed to the key child. For example, audio engine 208 (FIG. 3) can determine whether adult speech is directed to the key child or to someone else by determining the loudness/clearness of the adult speech or sounds associated with the segments. Once segmentation and segment ID are performed, segment-level near/far detection can be performed using the LRT in a manner similar to that used at the frame level. For each segment, the likelihood for each model can be calculated. The difference between the most probable model likelihood and the likelihood for the silence model can be calculated and the difference can be compared to a predetermined threshold. Based on the comparison, the segment can be either dropped or processed further.

**[0070]** Sometimes adults use baby talk or “parentese” when directing speech to children. The segments including parentese can be inaccurately associated with a child or the key child as the source because certain characteristics of the speech may be similar to speech-related vocalizations of the key child or other children. In many embodiments, audio engine 208 (FIG. 3) can modify the key child model and/or adult model to identify segments including parentese and associate the segments with an adult source. For example, the models may be modified to allow audio engine 208 (FIG. 3) to examine the complexity of the speech included in the segments to identify parentese. Because the complexity of adult speech is typically much higher than that of speech-related child vocalizations, the source for segments including relatively complex speech can be identified as an adult. Speech can be complex if the formant structures are well formed, the articulation levels are good, and the vocalizations are of sufficient duration – consistent with speech commonly provided by adults. Speech-related vocalization from a child can include formant structures that are less clear and developed and vocalizations that are typically of a lesser duration. In addition, audio engine 208 (FIG. 3) can analyze formant frequencies to identify segments including parentese. When an adult uses parentese, the formant frequencies of the segment typically do not change. Sources for segments including such identified parentese can be determined to be an adult.

**[0071]** The MD-GMM models can be further modified and the audio data can be further processed for a pre-set number of iterations or until audio engine 208 (FIG. 3) determines that the segment IDs have been determined with an acceptable level of confidence. Upon completion of the segmentation and segment ID, the identified segment can be further analyzed to extract characteristics associated with the language environment of the key child.

**[0072]** During or after performing segmentation and segment ID, audio engine 208 (FIG. 3) can classify key child audio segments into one or more categories, as described above. In some embodiments, audio engine 208 (FIG. 3) can analyze each segment for which the key child is identified as the source and can determine a category based on the sound in each segment. The categories can include vocalizations, cries, vegetative sounds, and fixed-signal sounds.

Vocalization sounds can include words, phrases, marginal syllables, including rudimentary consonant-vowel sequences, utterances, phonemes, sequence phonemes, phoneme-like sounds, protophones, lip-trilling sounds commonly called raspberries, canonical syllables, repetitive babbles, pitch variations, or any meaningful sounds which contribute to the language development of the child, indicate at least an attempt by the child to communicate verbally, or explore the capability to create sounds. Vegetative sounds can include non-vocal sounds related to respiration and digestion, such as coughing, sneezing, and burping. Fixed-signal sounds are related to voluntary reactions to the environment and include laughing, moaning, sighing, and lip smacking, and can be labeled as vegetative sounds. Cries are a type of fixed-signal sound but are detected separately. Cry segments can include any cry-related sound, such as intense cry (e.g., colicky cries), screams, wails, fussiness (e.g., non-colicky), and any other types of cry-related sounds.

**[0073]** In many embodiments, audio engine 208 (FIG. 3) may classify key child audio segments using rule-based analysis and/or statistical processing. Rule-based analysis can include analyzing each key child segment using one or more rules. For some rules, audio engine 208 may analyze energy levels or energy level transitions of segments. For example, a first rule that is based on a pre-set duration can be that segments including a burst of energy at or above the pre-set duration are identified as cries or screams and not vocalizations, but segments including bursts of energy less than the pre-set duration are classified as vocalizations. An example of one pre-set duration is three seconds based on characteristics commonly associated with vocalizations and cries. FIG. 5 illustrates energy levels of sound in a segment associated with the key child and showing a series of consonant (/b/) and vowel (/a/) sequences. Using a pre-set duration of three seconds, the bursts of energy indicate a vocalization since they are less than three seconds.

**[0074]** A second rule can be classifying segments as vocalizations that include formant transitions from consonant to vowel or vice versa. FIG. 5 illustrates formant transitions from consonant /b/ to vowel /a/ and then back to consonant /b/, indicative of canonical syllables and, thus, vocalizations. Segments that do not include such transitions may be further processed to determine a classification.

**[0075]** A third rule can be classifying segments as vocalizations if the formant bandwidth is narrower than a pre-set bandwidth. In some embodiments, the pre-set bandwidth is 1000 Hz based on common bandwidths associated with vocalizations.

**[0076]** A fourth rule can be classifying segments that include a burst of energy having a first spectral peak above a pre-set threshold as cries. In some embodiments, the pre-set threshold is 1500 Hz based on characteristics common in cries.

**[0077]** A fifth rule can be determining a slope of a spectral tilt and comparing it to pre-set thresholds. Often, vocalizations include more energy in lower frequencies, such as 300 to 3000 Hz, than higher frequencies, such as 6000 to 8000 Hz. A 30 dB drop is expected from the first part of the spectrum to the end of the spectrum, indicating a spectral tilt with a negative slope and a vocalization when compared to pre-set slope thresholds. Segments having a slope that is relatively flat may be classified as cries since the spectral tilt may not exist for cries. Segments having a positive slope can be classified as vegetative sounds.

**[0078]** A sixth rule can be comparing the entropy of the segment to entropy thresholds. Segments including relatively low entropy levels may be classified as vocalizations. Segments having high entropy levels can be classified as cries or vegetative sounds due to randomness of the energy.

**[0079]** A seventh rule can be comparing segment pitch to thresholds. Segments having a pitch between 250 to 600 Hz may be classified as vocalizations. Segments having a pitch of more than 600 Hz can be classified as cries or squeals, and a pitch of less than 250 Hz can be classified as growls.

**[0080]** An eighth rule can be determining pitch contours. Segments having a rising pitch may be classified as happy sounds. Segments having a falling pitch can be classified as angry sounds.

**[0081]** A ninth rule can be determining the presence of consonants and vowels. Segments having a mix of consonants and vowels may be classified as vocalizations. Segments having all or mostly consonants can be classified as vegetative sounds or fixed-signal sounds.

**[0082]** A rule according to various embodiments can be implemented separately or concurrently with other rules. For example, in some embodiments audio engine 208 (FIG. 3) can implement one rule only while, in other embodiments, audio engine 208 (FIG. 3) can implement two or more rules. Statistical processing can be performed in addition to or alternatively to the rule-based analysis.

**[0083]** Statistical processing can include processing segments with a MD-GMM using 2000 or more Gaussians in which models are created using Mel-scale Frequency Cepstral Coefficients (MFCC) and Spectral Subband Centroids (SSC). MFCCs can be extracted using a number of filter banks with coefficients. In one embodiment, 40 filter banks can be used with 36 coefficients. SSCs can be created using filter banks to capture formant peaks. The number of filter banks used to capture formant peaks can be 7 filter banks in the range of 300 to 7500 Hz. Other statistical processing can include using statistics associated with one or more of the following segment characteristics:

Formants;

Formant bandwidth;  
Pitch;  
Voicing percentage;  
Spectrum entropy;  
Maximum spectral energy in dB;  
Frequency of maximum spectral energy; and  
Spectral tilt.

Statistics regarding the segment characteristics can be added to the MFCC-SSC combinations to provide additional classification improvement.

**[0084]** As a child ages, characteristics associated with each key child segment category can change due to growth of the child's vocal tract. In some embodiments, an age-dependent model can be used in addition or alternatively to the techniques described above to classify key child segments. For example, vocalization, cry, and fixed-signal/vegetative-sound models can be created for each age group. In one embodiment, 12 different models are used with Group 1 corresponding to 1 to 2 months old, Group 2 corresponding to 3 to 4 months old, Group 3 corresponding to 5 to 6 months old, Group 4 corresponding to 7 to 8 months old, Group 5 corresponding to 9 to 10 months old, Group 6 corresponding to 11 to 12 months old, Group 7 corresponding to 13 to 14 months old, Group 8 corresponding to 15 to 18 months old, Group 9 corresponding to 19 to 22 months old, Group 10 corresponding to 23 to 26 months old, Group 11 corresponding to 27 to 30 months old, and Group 12 corresponding to 31 to 48 months old. In an alternative embodiment, vocalization, cry, and fixed-signal/vegetative-sound models may be created for each month of age from 1 month to 48 months. This model can include 144 models, 48 models for each category. Alternative embodiments can use a different number of groups or associate different age ranges with the groups.

**[0085]** In a number of embodiments, the cry-related segments can include any segment that is labeled as key child with the subcategory labeled as cry, and any segment that is labeled otherwise can be considered a non-cry segment. In other embodiments, other suitable segmenting and/or segment labeling approaches can be used. In many embodiments, the segments extracted from the datasets of audio data can include information about the start time of the segment, the end time of the segment, the duration (e.g., time length) of the segment, the average decibel (dB) sound pressure level of the segment, the peak dB sound pressure level of the segment, and/or other suitable information about the segment. In many embodiments, the average dB sound level of the segment can be determined by converting the dB values in the segment to sound power level (SPL) for averaging. For example, with certain recording or audio capture systems with an offset of 97, the conversion can be performed as follows:

$$\text{dB-PowerLevel} = 10^{((\text{Stored-dB}+97)/10)}$$

$$\text{dB-SPL} = \text{Round}(100 \cdot (10 \cdot \text{Log}_{10}(\text{dB-PowerLevel}))) / 100$$

**[0086]** Returning to FIG. 4, in several embodiments, method 400 additionally can include a block 415 of determining a plurality of periods of the plurality of cry-related segments that satisfy one or more threshold non-sparsity criteria. As an example, the determining of block 415 can be performed by device 200 (FIG. 3). In many embodiments, the plurality of periods can be separated by gaps of at least a first time-duration threshold without any cry-related segments. In some embodiments, the first time-duration threshold can be approximately 15 seconds to approximately 5 minutes. For example, in some embodiments, the first time-duration threshold can be approximately 30 seconds, such that the periods are separated by gaps of at least 30 seconds without any cry-related segments. In many embodiments, the period will start and end with a cry-related segment, which will be the same segment in the case of a period with a single cry-related segment. To illustrate, the first period can start with the first cry-related segment and will include any cry-related segments that occur within 30 seconds of each other without a gap of at least 30 seconds without a cry-related segment. Any gap of 30 seconds or more between cry-related segments will start a new period. In many embodiments, the cry-related segments are grouped into these periods sequentially. Often, the cry-related segments in the periods are non-continuous, and a period can have a number of cry-related segments as well as a number of non-cry-related segments. As an example, a period can include segments labeled in the following order: cry, cry, cry, noise (dog bark), cry, cry, female adult (mom soothing), cry, vegetation (intake of air by key child), cry, cry, cry, cry, female adult (mom soothing), cry, female adult (mom soothing), cry, other child, cry, other child, vegetation (intake of air by key child), cry, cry, cry, cry, noise (dog bark), cry, cry, female adult (mom soothing), cry, cry, vegetation (intake of air by key child), cry, silence, cry, cry.

**[0087]** Some of the periods can include only one or a limited number of cry-related segments, and such “sparse” periods can be filtered out by applying the one or more threshold non-sparsity criteria. In many embodiments, the periods that remain after applying the threshold non-sparsity criteria can be non-sparse periods of interest that are consequential.

**[0088]** In many embodiments, the one or more threshold non-sparsity criteria can include a first non-sparsity criterion in which each of the plurality of periods that satisfies the first non-sparsity criterion has a period duration longer than a second time-duration threshold. In some embodiments, the second time-duration threshold can be approximately 10 seconds to approximately 5 minutes. For example, in some embodiments, the second time-duration threshold can be set to 30 seconds, such that any period that is shorter than 30 seconds is filtered out. In many embodiments, this first non-sparsity criterion can advantageously ensure that the



period is long enough to have sufficient context data to be used in the classification performed in block 420, as described below in further detail.

**[0089]** In a number of embodiments, the one or more threshold non-sparsity criteria can include a second non-sparsity criterion in which each period of the plurality of periods that satisfies the second non-sparsity criterion has an aggregated cry duration that exceeds a third time-duration threshold. In many embodiments, the aggregated cry duration for each period can be an aggregate of durations of the plurality of cry-related segments within the period. In several embodiments, the third time-duration threshold can be approximately 5 seconds to approximately 30 seconds. For example, each of the cry-related segments in the period can have a duration, and these durations can be aggregated across the period to determine the aggregated cry duration of the period. In some examples, the third-time duration threshold can be 15 seconds, such that any periods in which the aggregated cry duration is less than 15 seconds are filtered out. In many embodiments, this second non-sparsity criterion can beneficially ensure that the period has enough cry segments to be considered non-sparse. In many embodiments, the third-time duration threshold can be adjusted based on the second-time duration threshold. In other embodiments, other threshold non-sparsity criteria can be used.

**[0090]** In several embodiments, the periods that are filtered out by the threshold non-sparsity criteria are not used in the classification performed in block 420, as described below in further detail. The periods that do satisfy the threshold non-sparsity criteria can be classified in block 420, as described below in further detail. Grouping cry-related segments into non-sparse periods can advantageously eliminate insubstantial cries and facilitate examining cry-related segments in a way that essentially averages the cry-related segments in a period to improve statistical analysis and performance. In many embodiments, a period can be a quasi- or reasonably-homogenous time sequence that can be further classified in block 420, as described below in further detail.

**[0091]** Turning ahead in the drawings, FIG. 6 illustrates a cry-related segment plot 600 showing a circular symbol for each cry-related segment that is labeled as such over a 720-minute (12-hour) audio recording, based on the elapsed capture time at which the cry-related segment occurs, and the resulting cumulative minutes of cry at which the cry-related segment occurs. As can be seen in FIG. 6, some regions of the cry-related segments are denser than others, and the slope or steepness of the staircase pattern varies over time.

**[0092]** Turning ahead to the next drawing, FIG. 7 illustrates a cry-related segment plot 700 in which the cry-related segments from cry-related segment plot 600 (FIG. 6) that are part of periods that satisfy the threshold non-sparsity criteria are labeled with a diamond symbol, and the other cry-related segments from cry-related segment plot 600 (FIG. 6) that are filtered out by

the non-sparsity criteria are labeled with a circular symbol.

**[0093]** Returning to FIG. 4, in a number of embodiments, method 400 further can include a block 420 of performing a classification on the plurality of periods to classify each of the plurality of periods as either a cry period or a fussiness period. In many embodiments, the classification can be a machine-learning algorithm that is based on labeled training data. For example, the training data can be periods that are labeled by expert listeners who have analyzed the periods in the training data. As an example, the performing of block 420 can be performed by device 200 (FIG. 3).

**[0094]** In some embodiments, one or more features (e.g., predictor variables) can be identified from the periods and/or other characteristics of the dataset of audio data or characteristics of the key child, which can be used in the machine learning classification algorithm. In many embodiments, one or more of the features can be period parameters. In various embodiments, various metrics related to the period can be calculated. For examples, the metrics can include the duration of the period, the aggregated cry duration of the period, the number of cry-related segments in the period, the mean and/or median duration of the cry-related segments in the period, the mean and/or median duration of the gaps of non-cry-related segments within the period, or other suitable metrics. One or more of these metrics can be used as a feature in the classification algorithm, or other suitable metrics can be used. For example, one of the features can be a slope of the period. The slope of the period can be based on the aggregate cry duration in the period relative to the total duration of the period. The slope of the period can be viewed as the steepness of the cry-related segments in the period in the cry-related segment plots (e.g., 600 (FIG. 6), 700 (FIG. 7)). As another example, another one of the features can be a median dB-SPL of the period. The median dB-SPL of the period can be the median of the average dB-SPL for all of the cry-related segments in the period. In other embodiments, other and/or additional period parameters can be used as features in the classification algorithm.

**[0095]** In a number of embodiments, one or more of the features can be audio dataset parameters. For example, the total number of periods across the dataset of audio data, including the periods that satisfy and the periods that do not satisfy the threshold non-sparsity criteria, can be a feature used in the classification algorithm. In other embodiments, other and/or additional suitable audio dataset parameters can be used as features in the classification algorithm.

**[0096]** In several embodiments, one or more of the features can be child parameters. For example, the gender of the key child and/or the age (e.g., days of age) of the key child can be features in the classification. In other embodiments, other and/or additional suitable child parameters can be features in the classification algorithm.

**[0097]** In some embodiments, performing the classification on the plurality of periods can

include applying a binary logistic regression model on features of the plurality of periods. For example, the binary logistic regression model can use one or more of the features (i.e., predictor variables) described above to classify each period as either a cry period or a fussiness period. In several embodiments, the features of the plurality of periods used in the binary logistic regression model can include the slope of the period, as described above, and the median decibel sound pressure level (dB-SPL) of the period, as described above. In many embodiments, the features also can include the total number of periods across the dataset of audio data, as described above, the gender of the key child, and the age of the key child. In other embodiments, other suitable features can be used in a binary logistic regression model. In other embodiments, another suitable machine learning classification algorithm can be used. For example, in some embodiments, a neural network classification can be used to classify each of the periods as either a cry period or a fussiness period. In other embodiments, a classification algorithm can assign other classifications to periods in addition to cry or fussiness, e.g., mixed cry and fussiness periods.

**[0098]** To evaluate this approach, a test was performed in which key child cry segments were extracted from 64 audio recordings of 41 children ages 81-120 days old (11-17 weeks old) to create a dataset of labeled training data (21 of the children had more than one recording). From this dataset were identified 138 audio clips of periods that satisfied the threshold non-sparsity criteria described above (i.e., first time-duration threshold = 30 seconds, second time-duration threshold = 30 seconds, third time-duration threshold = 15 seconds). 88% of the periods were determined to be sparse, and were filtered out, and 12% of the periods satisfied the threshold non-sparsity criteria. Yet these non-sparse periods accounted for 58% of the total duration of all the periods (both sparse and non-sparse), meaning that the non-sparse periods tended to be fairly short in duration. Also, the non-sparse periods included 75% of the total duration of all cry-related segments in the dataset, meaning that most of the cry-related segments were in the non-sparse periods. The human listener then identified homogeneous regions within these 138 non-sparse periods, and classified these regions within the periods as cry, fussiness, or other. The other category included sporadic cry as well as non-cry sounds, such as child vocalizing, mother talking, TV, dog whining, etc.). 108 of these clips were determined to be primarily (i.e., at least 55%) cry or fussiness. The 30 other clips were determined to be not sufficiently homogenous as cry or fussiness to be used in the training data. Specifically, 14 of these 30 other clips were determined to be primarily other (not primarily cry or primarily fussiness). The remaining 16 of these 30 other clips were determined to be not as homogenous as the other 108 clips.

**[0099]** Turning ahead in the drawings, FIG. 8 illustrates a stacked bar chart 800 showing the 108 training clips based on the percentage of the duration of the cry (in diagonal), fussiness (in

clear), and other (in solid) categories in each of the training clips. As explained above, the categorization of the regions of each of the clips was performed by human listeners. These clips are arranged in FIG. 8 such that the clips that are primarily cry are shown on the left, and the clips that are primarily fussiness are shown on the right. As seen in FIG. 8, in most of these clips, there are at least two of the categories, and in several clips, all three categories are found.

**[0100]** Using the five features described above (the slope of the period, the median decibel sound pressure level (dB-SPL) of the period, the total number of periods across the recording, the gender of the key child, and the age of the key child, a binary logistic regression model was applied to each human-labeled training data period (108 clips) to generate log odds that the period is a cry period. Based on the training data, the binary logistic regression model output the log odds classifier shown in Equation 1:

$$\text{Log Odds} = -45.4640 + 20.0381 * \text{Slope} + 0.7785 * \text{Median Average dB SPL} - 4.9238 * (\text{Gender} = \text{Male}) - 3.7381 * \text{Child Age} - 0.0975 * \text{Total Periods} \quad (\text{Eq. 1})$$

The probability that the period is a cry period can then be determined from the Log Odds as shown in Equation 2:

$$\text{Cry Probability} = \text{EXP}(\text{Log Odds}) / (1 + \text{EXP}(\text{Log Odds})) \quad (\text{Eq. 2})$$

**[0101]** To test this classifier, the classifier was applied to the 138 human-labeled clips, and the threshold probability was set at just below 0.5, such that those periods with a probability above the threshold probability were labeled as cry periods, and those periods with a probability below the threshold probability were labeled as fussiness periods. As described earlier, 14 of the 138 human-labeled clips (or 10%) were primarily other and not primarily cry or primarily fussiness. These are False Positive, or Type I errors. These errors generally resulted from segment labeling errors, such as dog whining or electric toothbrush hum being mistaken for key child cry. Thirteen of these 14 clips were classified as fussiness periods, and the other clip was classified as a cry period using this classifier model, indicating that it was rare for the classifier to falsely identify a period of interest as a cry period.

**[0102]** For the remaining 124 human-labeled clips, the confusion matrix for the classifier is shown below in Table 1, which indicates balanced performance. Specifically, the sensitivity, meaning the probability of the classifier predicting cry when it is truly cry was 90%. The specificity, meaning the probability of the classifier predicting fuss when it is truly fuss was 92%. The positive predictive power (PPP), meaning the number of true cry periods versus the number of cry periods so called by the classifier was 92%. The negative predictive power (NPP), meaning the number of true fuss periods versus the number of fuss periods so called by the algorithm was 91%.

TABLE 1

	Predicted Cry	Predicted Fuss	Total	
Observed Cry	55	6	61	90% Sensitivity
Observed Fuss	5	58	63	92% Specificity
Total	60	64	124	
	92% PPP	91% NPP		91% Overall Accuracy

**[0103]** Turning ahead to the next drawing, FIG. 9 illustrates a probability chart 900 showing the predicted value for cry probability across the 124 human-labeled clips, separated into two rows, on an expanded probability axis. The top row shows the predicted values for the 63 clips that were labeled as fuss by the human listener, and the bottom row shows the predicted values for the 61 clips that were labeled as cry by the human listener. The decision threshold was set to slightly less than 0.5, which was done to achieve a balanced (quasi-equal) error rate.

**[0104]** Turning ahead to the next drawing, FIG. 10 illustrates a bubble frequency chart 1000 showing the predicted value for cry probability across the 124 human-labeled clips, separated into two rows, on an expanded probability axis, using bubble frequency. FIG. 10 provides an alternate way of viewing the data in FIG. 9. Like FIG. 9, the top row shows the predicted values for the 63 clips that were labeled as fuss by the human listener, and the bottom row shows the predicted values for the 61 clips that were labeled as cry by the human listener. As shown in FIG. 10, the clips are grouped at various probability levels with the size of the circle corresponding to the frequency of the clips that fall within the probability level. FIG. 10 shows that the classifier provided good separation between the two classifications of fussiness and cry, as the majority of the cases are at or near the two extreme ends of the probability axis.

**[0105]** Turning ahead to the next drawing, FIG. 11 illustrates a cry-related segment plot 1100 in which the cry-related segments that are part of periods that satisfy the threshold non-sparsity criteria in FIG. 7 (shown as the diamond symbols in FIG. 7), are classified using the binary logistic regression model described above as either part of a cry period (and labeled with diamond symbols) or as part of a fussiness period (and labeled with square symbols). In other words, each cry-related segment in a period is coded with a symbol according to the classification of the period of which it is a part. Based on the features used in the classifier, general trends show that, on the one hand, cry probability increases for steeper slope, higher median dB-SPL, fewer total periods, gender as male, and younger age. On the other hand, general trends show that fussiness probability increases for gentler slope, lower median dB-SPL, more total periods, gender as female, and older age.

**[0106]** Returning in the drawings to FIG. 4, in several embodiments, method 400 optionally can

include a block 425 of outputting metrics based on the classification of the plurality of periods. Metrics can include any quantifiable value that uses the classification of the periods. For example, the metrics can include the number of periods classified as cry periods and/or fussiness periods, the total duration of the periods classified as cry periods and/or fussiness periods, the total duration of the cry-related segments in the periods classified as cry periods and/or fussiness periods, the number of cry-related segments in segments in the periods classified as cry periods and/or fussiness periods, and/or other suitable metrics. In addition, the output can include information such as the date and time of the collection of the dataset of audio data, the duration of the dataset of audio data, the age and/or gender of the key child, period-level details about the features, and/or other suitable metrics.

**[0107]** As an example, the outputting of block 425 can be performed by device 200 (FIG. 3). In many embodiments, audio engine 208 (FIG. 3) can output the metrics to output device 214 (FIG. 3). For example, audio engine 208 (FIG. 3) can, in response to a command received from input device 212 (FIG. 3), output a metric associated with the classification of the periods and/or the cry-related segments. In many embodiments, the systems and methods described herein can beneficially provide objective consistency in quantifying the amount of cry from a key child, which can be used to diagnose infantile colic more accurately and/or to assess whether an intervention (e.g., medicine, a change in diet, etc.) is effective at mitigating infantile colic.

**[0108]** Although automatic determination of cry segments and discrimination of cry from fussiness has been described with reference to specific embodiments, it will be understood by those skilled in the art that various changes may be made without departing from the spirit or scope of the disclosure. Accordingly, the disclosure of embodiments is intended to be illustrative of the scope of the disclosure and is not intended to be limiting. It is intended that the scope of the disclosure shall be limited only to the extent required by the appended claims. For example, to one of ordinary skill in the art, it will be readily apparent that any element of FIGs. 1-11 may be modified, and that the foregoing discussion of certain of these embodiments does not necessarily represent a complete description of all possible embodiments. For example, one or more of the procedures, processes, or activities of FIG. 4 may include different procedures, processes, and/or activities and be performed by many different modules, in many different orders.

**[0109]** Replacement of one or more claimed elements constitutes reconstruction and not repair. Additionally, benefits, other advantages, and solutions to problems have been described with regard to specific embodiments. The benefits, advantages, solutions to problems, and any element or elements that may cause any benefit, advantage, or solution to occur or become more pronounced, however, are not to be construed as critical, required, or essential features or

elements of any or all of the claims, unless such benefits, advantages, solutions, or elements are stated in such claim.

**[0110]** Moreover, embodiments and limitations disclosed herein are not dedicated to the public under the doctrine of dedication if the embodiments and/or limitations: (1) are not expressly claimed in the claims; and (2) are or are potentially equivalents of express elements and/or limitations in the claims under the doctrine of equivalents.

## CLAIMS

What is claimed is:

1. A system comprising:
  - one or more processors; and
  - one or more non-transitory computer-readable media storing computing instructions configured to run on the one or more processors and perform:
    - receiving one or more datasets of audio data of a key child captured in a natural sound environment of the key child;
    - segmenting each of the one or more datasets of audio data to create audio segments, the audio segments comprising cry-related segments and non-cry segments;
    - determining periods of the cry-related segments that satisfy one or more threshold non-sparsity criteria; and
    - performing a classification on the periods to classify each of the periods as either a cry period or a fussiness period.
2. The system of claim 1, wherein the computing instructions are further configured to perform:
  - outputting metrics based on the classification of the periods.
3. The system of any one of claims 1 or 2, wherein the periods are separated by gaps of at least a first time-duration threshold without any cry-related segments.
4. The system of claim 3, wherein the first time-duration threshold is approximately 15 seconds to approximately 5 minutes.
5. The system of any one of claims 1, 2, 3, or 4, wherein the one or more threshold non-sparsity criteria comprise a first non-sparsity criterion in which each of the periods that satisfies the first non-sparsity criterion has a period duration longer than a second time-duration threshold.
6. The system of claim 5, wherein the second time-duration threshold is approximately 10 seconds to approximately 5 minutes.
7. The system of any one of claims 1, 2, 3, 4, 5, or 6, wherein:



the one or more threshold non-sparsity criteria comprise a second non-sparsity criterion in which each period of the periods that satisfies the second non-sparsity criterion has an aggregated cry duration that exceeds a third time-duration threshold; and the aggregated cry duration for the each period is an aggregate of durations of the cry-related segments within the each period.

8. The system of claim 7, wherein the third time-duration threshold is approximately 5 seconds to approximately 30 seconds.

9. The system of any one of claims 1, 2, 3, 4, 5, 6, 7, or 8, wherein the one or more datasets of audio data of the key child comprise datasets of audio data in which the key child is less than 4 months old.

10. The system of any one of claims 1, 2, 3, 4, 5, 6, 7, 8, or 9, wherein performing the classification on the periods comprises:

applying a binary logistic regression model on features of the periods.

11. The system of claim 10, wherein:

the features of the periods used in the binary logistic regression model comprise:

a slope of the period;

a median decibel sound pressure level of the period;

a total number of the periods of the cry-related segments, including (a) the periods of the cry-related segments that satisfy the one or more threshold non-sparsity criteria and (b) the periods of the cry-related segments that do not satisfy the one or more threshold non-sparsity criteria;

an age of the key child; and

a gender of the key child.

12. A method being implemented via execution of computing instructions configured to run on one or more processors and stored on one or more non-transitory computer-readable media, the method comprising:

receiving one or more datasets of audio data of a key child captured in a natural sound environment of the key child;  
segmenting each of the one or more datasets of audio data to create audio segments, the audio segments comprising cry-related segments and non-cry segments;  
determining periods of the cry-related segments that satisfy one or more threshold non-sparsity criteria; and  
performing a classification on the periods to classify each of the periods as either a cry period or a fussiness period.

13. The method of claim 12 further comprising:

outputting metrics based on the classification of the periods.

14. The method of any one of claims 12 or 13, wherein the periods are separated by gaps of at least a first time-duration threshold without any cry-related segments.

15. The method of claim 14, wherein the first time-duration threshold is approximately 15 seconds to approximately 5 minutes.

16. The method of any one of claims 12, 13, 14, or 15, wherein the one or more threshold non-sparsity criteria comprise a first non-sparsity criterion in which each of the periods that satisfies the first non-sparsity criterion has a period duration longer than a second time-duration threshold.

17. The method of claim 16, wherein the second time-duration threshold is approximately 10 seconds to approximately 5 minutes.

18. The method of any one of claims 12, 13, 14, 15, 16, or 17, wherein:

the one or more threshold non-sparsity criteria comprise a second non-sparsity criterion in which each period of the periods that satisfies the second non-sparsity criterion has an aggregated cry duration that exceeds a third time-duration threshold; and  
the aggregated cry duration for the each period is an aggregate of durations of the cry-related segments within the each period.

19. The method of claim 18, wherein the third time-duration threshold is approximately 5 seconds to approximately 30 seconds.
20. The method of any one of claims 12, 13, 14, 15, 16, 17, 18, or 19, wherein the one or more datasets of audio data of the key child comprise datasets of audio data in which the key child is less than 4 months old.
21. The method of any one of claims 12, 13, 14, 15, 16, 17, 18, 19, or 20, wherein performing the classification on the periods comprises:  
    applying a binary logistic regression model on features of the periods.
22. The method of claim 21, wherein:  
    the features of the periods used in the binary logistic regression model comprise:  
        a slope of the period;  
        a median decibel sound pressure level of the period;  
        a total number of the periods of the cry-related segments, including (a) the periods of the cry-related segments that satisfy the one or more threshold non-sparsity criteria and (b) the periods of the cry-related segments that do not satisfy the one or more threshold non-sparsity criteria;  
        an age of the key child; and  
        a gender of the key child.

1/11






FIG. 1

2/11



FRONT VIEW

FIG. 2A



SIDE VIEW

FIG. 2B

3/11

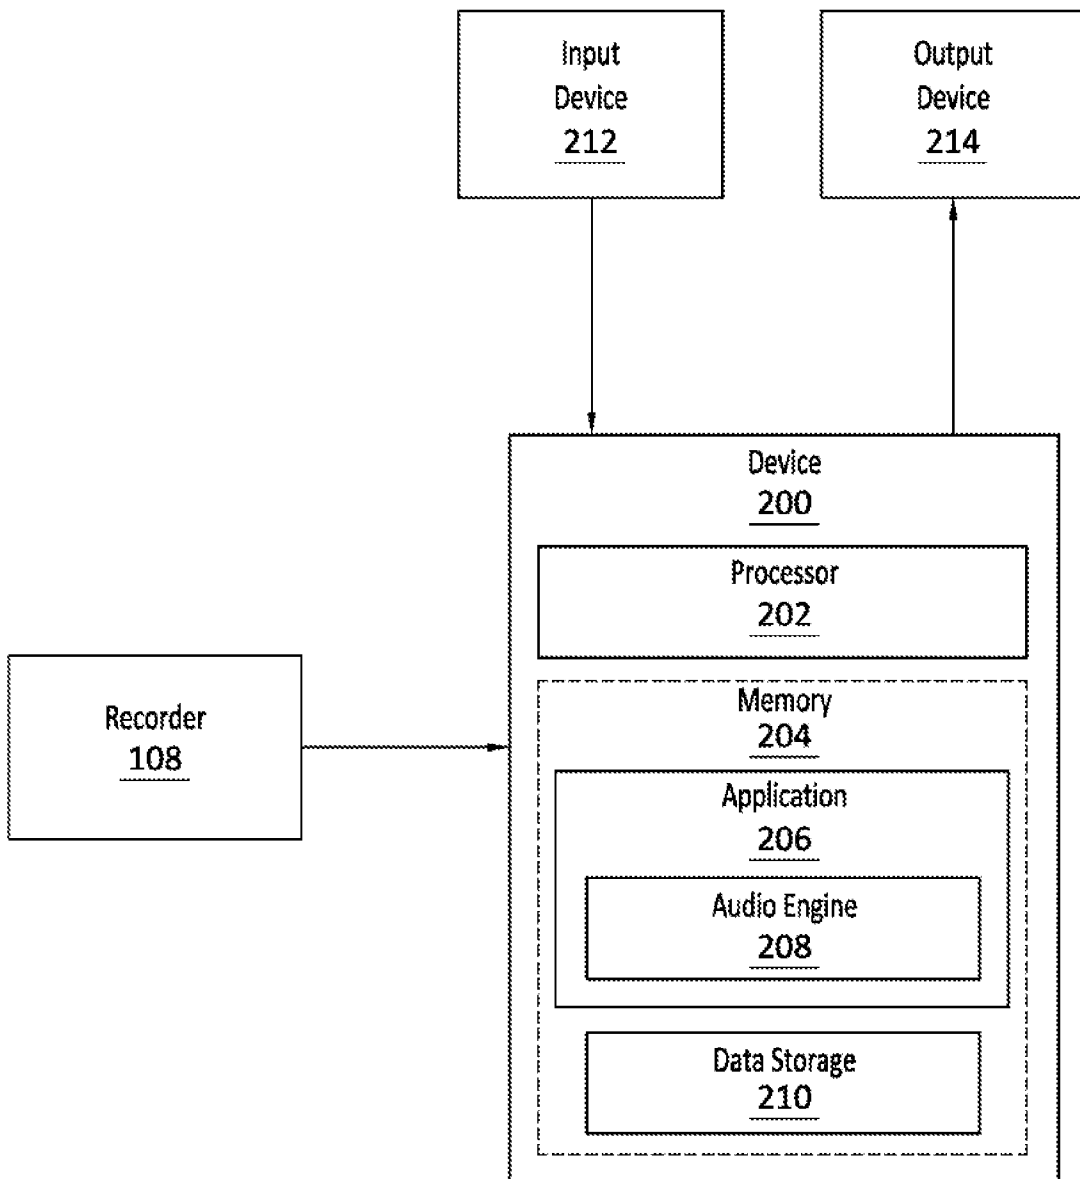


FIG. 3

4/11

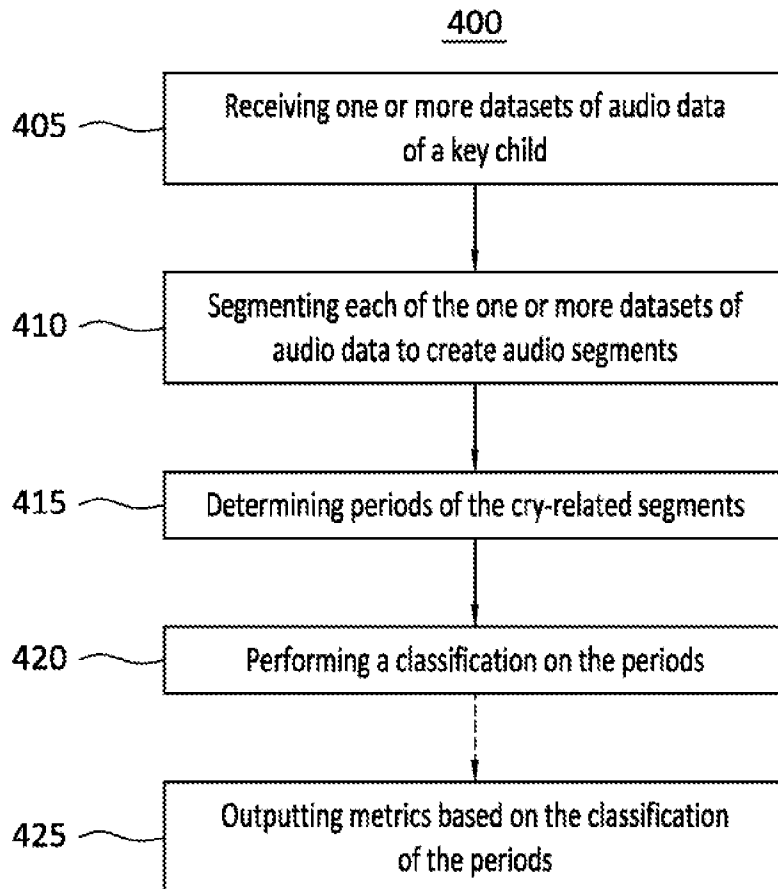


FIG. 4




FIG. 5



6/11




FIG. 6

7/11




FIG. 7

8/11




FIG. 8

9/11

900

Predicted Value For Cry Probability





FIG. 9



**FIG. 10**

11/11




FIG. 11

## INTERNATIONAL SEARCH REPORT

International application No.  
**PCT/US2018/064523****A. CLASSIFICATION OF SUBJECT MATTER****G10L 17/26(2013.01)i, G10L 25/78(2013.01)i**

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 17/26; G06F 17/20; G10L 15/16; G10L 21/0216; G10L 25/30; G10L 25/48; G10L 25/78Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched  
Korean utility models and applications for utility models  
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)  
eKOMPASS(KIPO internal) & Keywords: audio data, key child, cry-related segments, classification, gaps, periods**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	WO 2017-112261 A1 (INTEL IP CORPORATION) 29 June 2017 See paragraphs [0018], [0032], [0040]-[0045]; claim 10; and figures 2, 6A-6B, 8.	1-22
Y	US 9009038 B2 (NATIONAL TAIWAN NORMAL UNIVERSITY) 14 April 2015 See column 3, lines 19-22; column 5, lines 34-35; claim 1; and figure 1.	1-22
Y	YIZHAR LAVNER et al., 'Baby Cry Detection in Domestic Environment using Deep Learning', 2016 IEEE International Conference on the Science of Electrical Engineering, 16 November 2016 See section III; and figure 3.	10-11, 21-22
A	LINA ABOU-ABBAS et al., 'A fully automated approach for baby cry signal segmentation and boundary detection of expiratory and inspiratory episodes', The Journal of the Acoustical Society of America, Volume 142, Issue 3, 11 September 2017, pages 1318-1331. See pages 1325-1327; and figure 13.	1-22
A	US 2015-0112678 A1 (DOMINIC FRANK JULIAN BINKS et al.) 23 April 2015 See paragraphs [0049]-[0057]; claim 1; and figure 1a.	1-22

 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

"&amp;" document member of the same patent family

Date of the actual completion of the international search

27 March 2019 (27.03.2019)

Date of mailing of the international search report

**27 March 2019 (27.03.2019)**

Name and mailing address of the ISA/KR

International Application Division  
Korean Intellectual Property Office  
189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

YANG, Jeong Rok

Telephone No. +82-42-481-5709



**INTERNATIONAL SEARCH REPORT**

Information on patent family members

International application No.

**PCT/US2018/064523**

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2017-112261 A1	29/06/2017	US 2017-0178667 A1 US 9899034 B2	22/06/2017 20/02/2018
US 9009038 B2	14/04/2015	CN 103426438 A CN 103426438 B TW 201349224 A TW I474315 B US 2013-0317815 A1	04/12/2013 09/12/2015 01/12/2013 21/02/2015 28/11/2013
US 2015-0112678 A1	23/04/2015	CN 102246228 A CN 102246228 B GB 2466242 A GB 2466242 B GB 2534027 A GB 2534027 A8 GB 2534027 B GB 2534027 B8 US 2011-0218952 A1 US 2015-0106095 A1 US 8918343 B2 US 9286911 B2 WO 2010-070314 A1	16/11/2011 18/06/2014 23/06/2010 02/01/2013 13/07/2016 06/12/2017 19/07/2017 06/12/2017 08/09/2011 16/04/2015 23/12/2014 15/03/2016 24/06/2010