

CLAIMS:

1. A method for identifying segments of a digital audio recording of sounds from a subject, where the segments contain particular sound events of interest, the method comprising:

filtering the digital audio recording based on a characteristic frequency range of the sound events to produce a filtered digital audio signal;

processing the filtered digital audio signal to produce a corresponding signal envelope;

fitting a statistical distribution to the signal envelope;

determining a threshold level for the signal envelope based on the statistical distribution and a predetermined probability level; and

identifying segments of the signal envelope that are above the threshold level to thereby identify corresponding segments of the digital audio recording of sounds from the subject as segments of the digital audio recording containing the particular sound events of interest.

2. The method of claim 1, including applying a first downsampling by which a sample rate of the digital audio recording is reduced by an integer factor to produce a first downsampled digital audio signal.

3. The method of claim 2, wherein the first downsampled digital audio signal is filtered in the characteristic frequency range to select for the sound events of interest to thereby produce a first downsampled and event-filtered digital audio signal.

4. The method of claim 3, wherein the events of interest comprise breath sounds and wherein filtering the digital audio recording comprises applying a high pass filter.

5. The method of claim 3, wherein the events of interest comprise snore sounds and wherein filtering the digital audio recording comprises applying a low pass filter.
6. The method of any one of claim 3 to claim 5, wherein processing the filtered digital audio signal to produce a corresponding signal envelope is implemented by an envelope detection procedure.
7. The method of claim 6, wherein the envelope detection procedure includes applying an absolute value filter to the first downsampled and event-filtered signal to produce an absolute value filtered signal.
8. The method of claim 7, wherein the absolute value filtered signal is filtered by a forward and reverse filter to produce a low pass filtered absolute value signal.
9. The method of claim 8, including applying a second downsampling to the low pass filtered absolute value signal to produce the signal envelope, the signal envelope comprising a first signal envelope which is an estimate of amplitude of the audio recording.
10. The method of claim 9, wherein applying the second downsampling comprises resampling from 14.7kHz down to 100Hz.
11. The method of claim 9 or claim 10, including applying logarithmic compression to the first signal envelope to produce a second signal envelope that comprises a power estimate of the digital audio recording.
12. The method of claim 11, wherein fitting the statistical distribution to the signal envelope comprises fitting the statistical distribution to the second signal envelope that comprises the power estimate.

13. The method of any one of the preceding claims, wherein fitting the statistical distribution to the signal envelope includes sorting samples making up the signal envelope into a number of bins to produce a histogram.

14. The method of claim 13, wherein fitting the statistical distribution includes selecting a modal bin of the histogram, wherein the modal bin is a bin into which the greatest number of samples have been sorted.

15. The method of claim 14, wherein the statistical distribution comprises a Poisson distribution having a lambda parameter and fitting the statistical distribution includes setting the lambda parameter to the number of the modal bin.

16. The method of any one of the preceding claims, wherein determining the threshold level for the signal envelope based on the statistical distribution and the predetermined probability level comprises calculating a cumulative distribution function (CDF) in respect of the statistical distribution.

17. The method of claim 16, wherein determining the threshold level for the signal envelope comprises finding a threshold bin, being a bin that corresponds to the predetermined probability, wherein the predetermined probability level comprises a probability level on the CDF.

18. The method of claim 17, wherein determining the threshold level comprises setting the threshold level to a value from a range of magnitudes of samples in the threshold bin.

19. The method of claim 17 or claim 18, wherein the threshold level is set to an upper limit of the range of magnitude of samples in the threshold bin.

20. The method of any one of claims 16 to 19, wherein a temporal filter is applied to cull segments that do not fall within a predetermined range of durations based on the events of interest.

21. The method of claim 20, wherein the events of interest comprise snore sounds and wherein the range of durations is greater than 225 milliseconds and less than 4 seconds.

22. The method of any one of the preceding claims, including recording information indicating start and end times for each of the segments of the signal envelope that are above the threshold level in a non-volatile manner and in association with the digital audio recording.

23. An apparatus comprising a sound event identification machine configured to identify portions of a digital audio recording of a subject containing a particular sound events of interest, including:

- a processor for processing the digital recording in accordance

- a digital memory in data communication with the processor, the digital memory storing instructions to configure the processor, the instructions including instructions configuring the processor to:

- filter the recording based on a characteristic frequency range of the sound events;

- process the filtered recording to produce a corresponding signal envelope;

- fit a statistical distribution to the signal envelope to thereby determine a threshold level corresponding to a predetermined probability level; and

- identify segments of the signal envelope that are above the threshold to thereby identify corresponding segments of the digital audio recording as segments containing the particular sound events.

24. The apparatus of claim 23, including a microphone that is configured to pick up sounds of the subject.

25. The apparatus of claim 23 or claim 24, including an audio interface comprising a filter and an analog-to-digital converter configured to convert the sounds of the subject into a digital audio signal.

26. The apparatus of claim 25, wherein the apparatus is configured to store the digital audio signal as the digital audio recording in the digital memory accessible to the processor.

27. The apparatus of any one of claims 23 to 26, including a human-machine-interface.

28. The apparatus of claim 27, wherein the instructions stored in the digital memory include instructions that configure the processor to display information on the human-machine-interface including information identifying segments in the digital audio recording containing the events of interest.

29. The apparatus of claim 28, wherein the instructions stored in the digital memory include instructions that configure the processor to display information on the human-machine-interface including information indicating the event of interest.

30. The apparatus of claim 28 or claim 29, wherein the information that is displayed on the human-machine-interface includes information indicating a start time and an end time in respect of each of a number of segments identified to contain the event of interest.

31. The apparatus of claim 30, wherein the digital memory includes instructions that configure the processor to write the start and end times for each identified segment in a non-volatile manner to thereby tangibly label segments containing the events of interest in respect of the digital audio recording.

32. A machine-readable media bearing tangible, non-transitory instructions for execution by one or more processors to implement the method of claim 1.

* * *