



(19) **United States**

(12) **Patent Application Publication**
Battaglino et al.

(10) **Pub. No.: US 2016/033553 A1**
(43) **Pub. Date: Nov. 17, 2016**

(54) **ACCOUSTIC CONTEXT RECOGNITION USING LOCAL BINARY PATTERN METHOD AND APPARATUS**

G06N 99/00 (2006.01)
G10L 21/10 (2006.01)

(52) **U.S. Cl.**

CPC *G06N 5/047* (2013.01); *G10L 21/10* (2013.01); *G10L 19/02* (2013.01); *G06N 99/005* (2013.01); *G10L 2019/0004* (2013.01)

(71) Applicant: **NXP B.V.**, Eindhoven (NL)

(72) Inventors: **Daniele Battaglino**, Alpes-Maritimes (FR); **Ludovick Lepauloux**, Cannes (FR); **Laurent Pilati**, Biot (FR); **Nicholas Evans**, Vallbonne (FR)

(57) **ABSTRACT**

(21) Appl. No.: **15/141,942**

(22) Filed: **Apr. 29, 2016**

(30) **Foreign Application Priority Data**

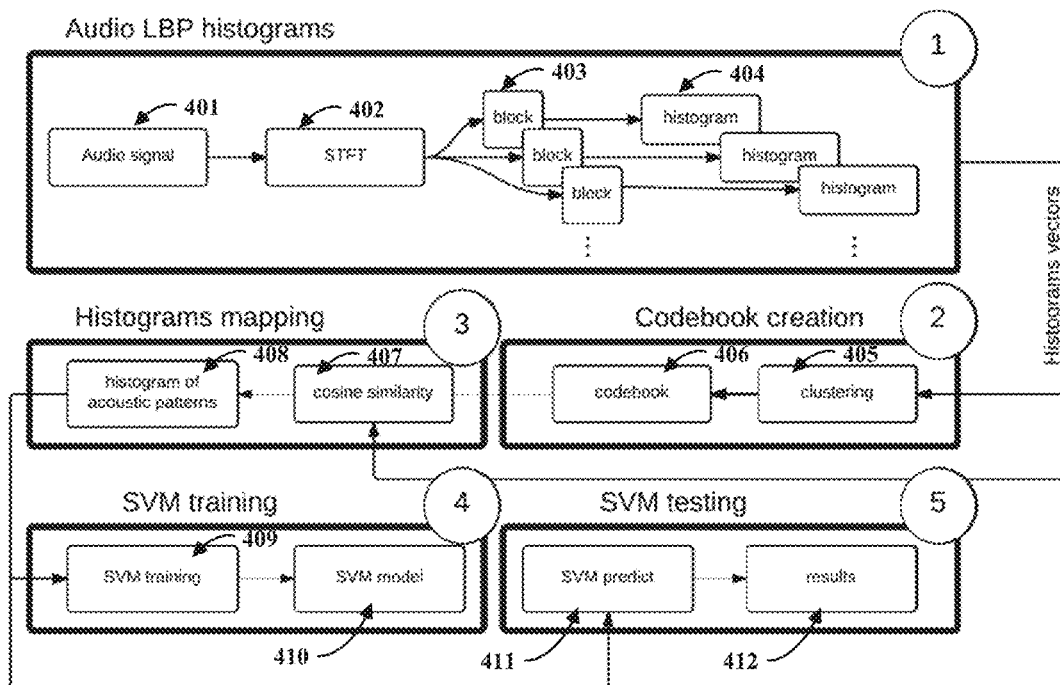
May 12, 2015 (EP) 15290128.6

Publication Classification

(51) **Int. Cl.**

G06N 5/04 (2006.01)
G10L 19/02 (2006.01)

Various exemplary aspects are directed to acoustic context recognition apparatuses and methods involving isolating and identifying context(s) of an acoustic environment. In one exemplary embodiment, source audio is converted into audio spectrograms, each spectrogram indicative of a period of time. The series of spectrograms are analyzed to identify audio patterns, over a period of time, which are indicative of an environmental context of the source audio. In many embodiments of the present disclosure, acoustic context recognition also includes comparing the identified audio patterns to known environmental contexts.



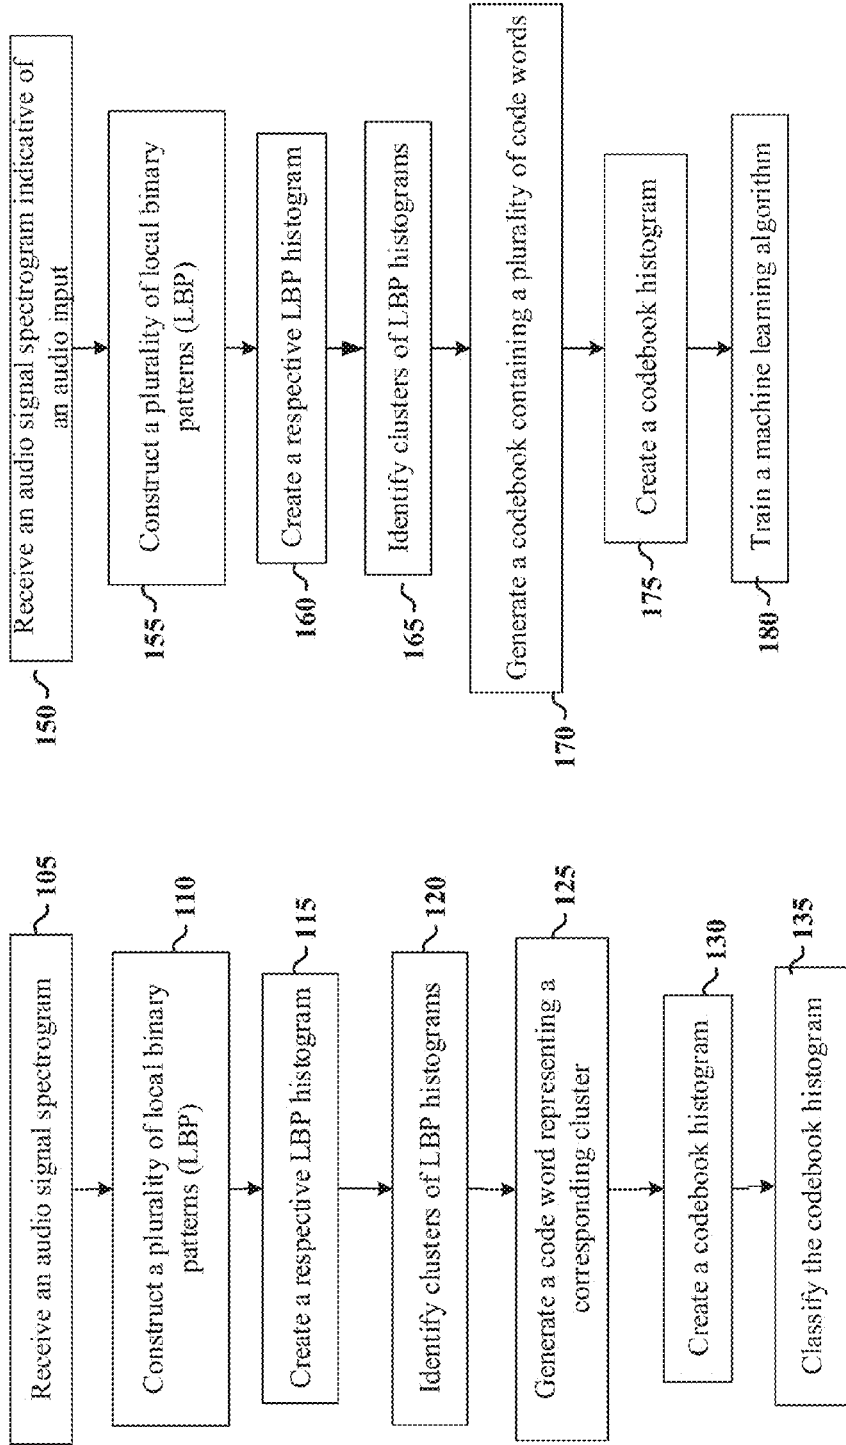


FIG. 1A

FIG. 1B

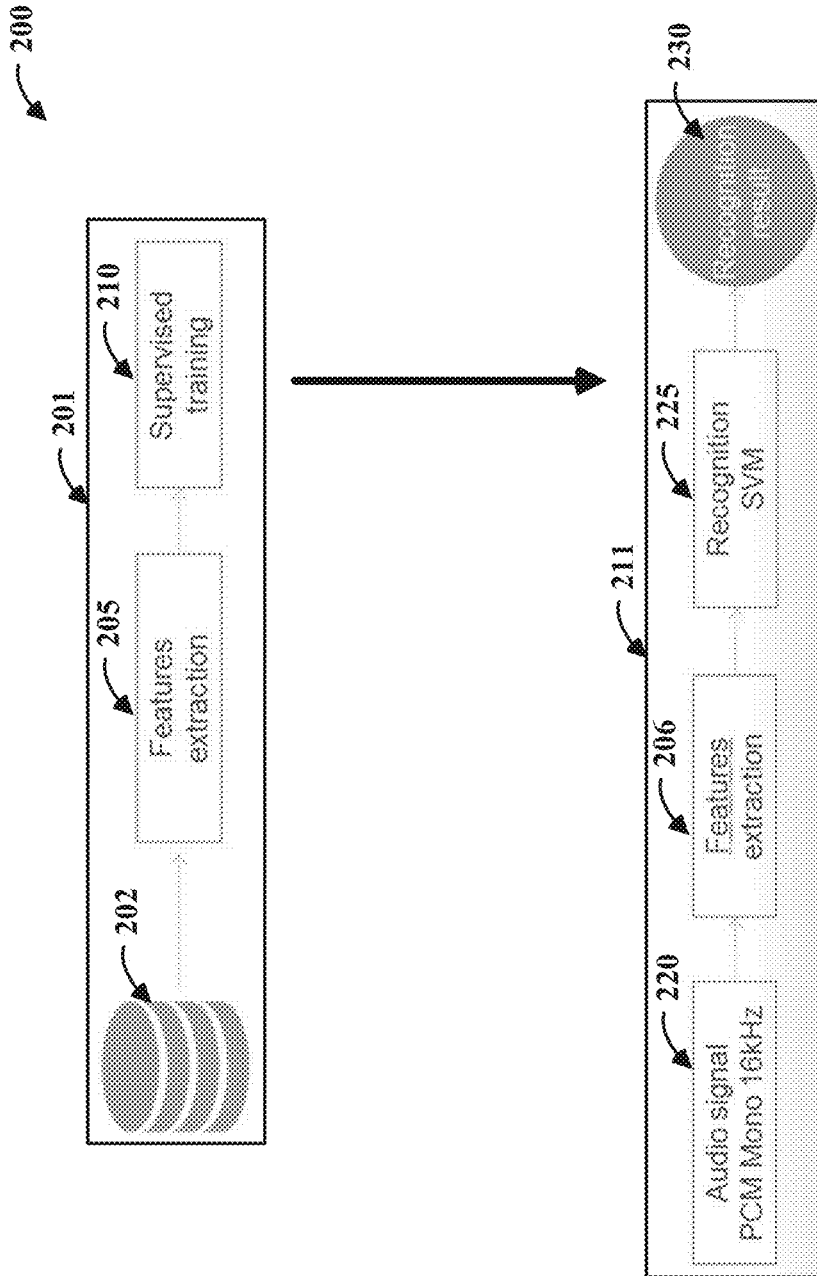


FIG. 2

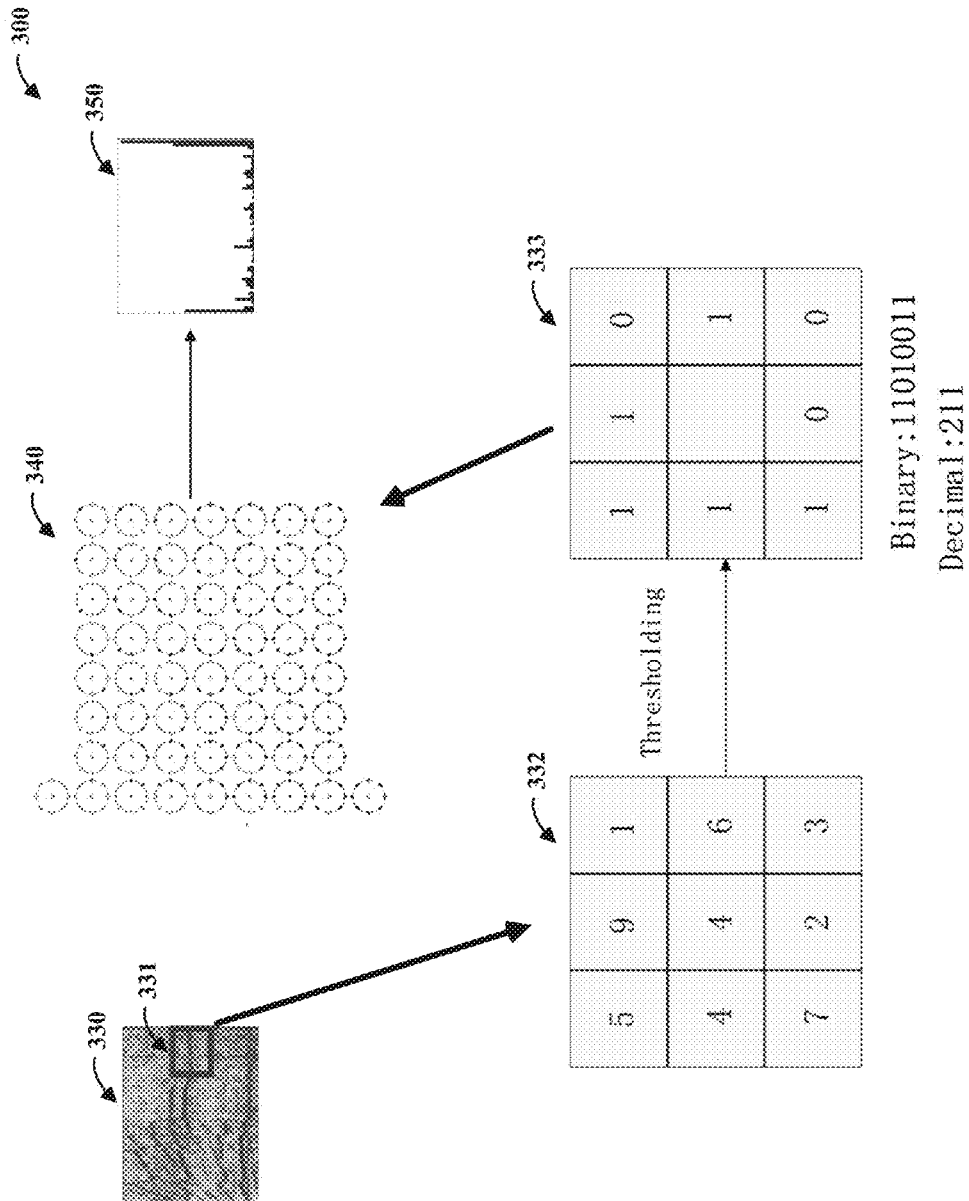


FIG. 3

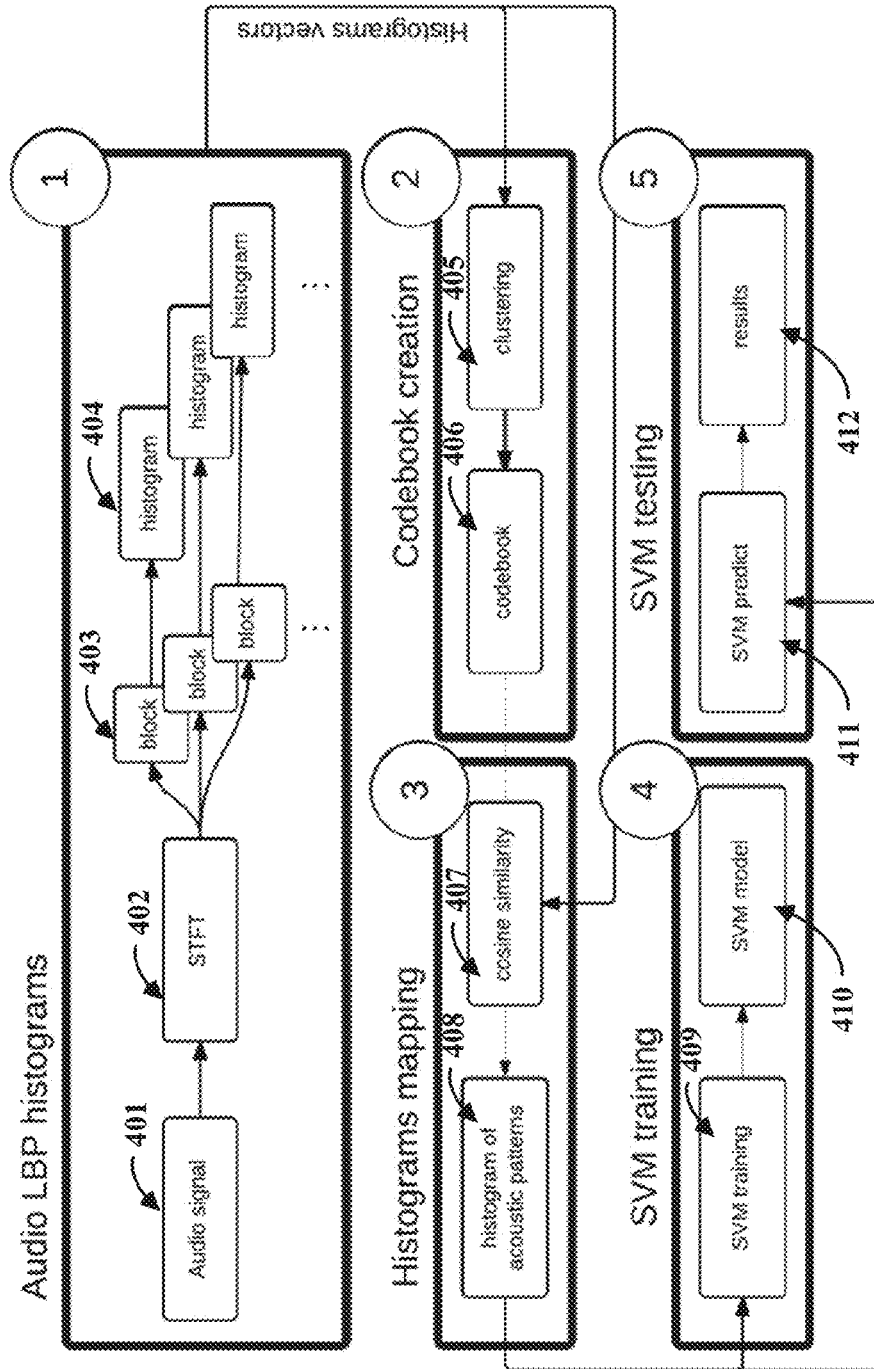


FIG. 4

**ACCOUSTIC CONTEXT RECOGNITION
USING LOCAL BINARY PATTERN METHOD
AND APPARATUS**

[0001] Aspects of the present disclosure generally relate to sound processing. In particular, the present disclosure presents solutions for audio scene recognition utilizing local binary patterns to identify patterns in audio which can be associated with, e.g., different sources, speech, music, background noise and particular events.

[0002] In many applications of computer electronics, (e.g., cell phones), customers require high levels of personalization, including the data that is presented to a user. Embodiments of the present disclosure are directed at providing such personalization by contextualizing the environment the device is being used in using audio received over a microphone. In such embodiments, the device is able to isolate and identify context(s) of the environment and present information to the user associated at least in part with the environment. The context of an audible environment may include various audio sources including sounds generally associated with a location, such as an office, bus, or street. The use of audio scene recognition over other methodologies provides several benefits including at least the ability to capture temporal events, and the ability to distinguish several co-occurring contexts in the environment. Various embodiments can take advantage of hardware present in electronic devices, such as microphones and sufficient processing circuitry. It has been discovered in accordance with the present disclosure that the ability to capture temporal events greatly increases the probability that a context will be properly identified by the device. In further more specific embodiments of the present disclosure, the device may further reduce audio processing requirements of the device, thereby increasing battery life if the device is mobile and improving overall performance of the device as less processing power needs to be allotted to background functions.

[0003] Aspects of the present disclosure apply image processing techniques on an audio spectrum by transforming the audio into an audio spectrogram (a visual representation of the audio at one or more periods of time) and analyzing a series of spectrograms from the same audio source to identify audio patterns that are indicative of an environmental context that a device receiving the audio is in (e.g., office, bus, street, etc.). The present disclosure's ability to map these audio patterns over time, using histograms associated with the series of spectrograms, greatly increases the acoustic recognition performance.

[0004] In many embodiments of the present disclosure, the acoustic recognition method may utilize a pre-defined codebook, whereby the identified audio patterns are compared to known environmental contexts. When compared, if the identified audio pattern is within a similarity threshold to one or more of the known audio patterns in the codebook, the identified audio pattern will be associated with the known audio pattern's environmental context. In such embodiments, the identified audio pattern may be added to the codebook and associated therein with the appropriate environmental context. In yet further embodiments, the method may include a learning-phase which allows a user of a device utilizing the acoustic recognition method to identify audio patterns (which the acoustic recognition method may not be able to identify) and add them to the codebook.

[0005] Various embodiments of the present disclosure are directed to identifying context(s) of an audio scene. Accord-

ing to one such example embodiment, a method is disclosed including the following steps. An audio signal spectrogram, indicative of an audio input, is received, and based on a comparison of adjacent pixels of the spectrogram, a plurality of local binary patterns (LBP) are constructed. For each of the plurality of blocks of the spectrogram, a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks is created. Based upon correspondence between the LBP histograms and a plurality of code words, a codebook histogram is created. Finally, using a machine learning model, the codebook histogram is classified to identify context for the received audio signal. In yet further more specific embodiments of the present disclosure, the method may further include identifying clusters of LBP histograms, with a k-means algorithm, from the respective LBP histograms, and generating, based upon the clusters, code words which are a centroid of the identified clusters. In certain embodiments, the classification and context identification can be carried out in real-time, or in near real-time.

[0006] Other embodiments of the present disclosure are directed to an offline learning phase for identifying context (s) of an audio scene (these identified contexts can later be used in real-time applications to identify contexts of a received audio input). According to one such example embodiment, a method is disclosed including the following steps. An audio signal spectrogram, indicative of an audio input, is received, and based on a comparison of adjacent pixels of the spectrogram, a plurality of local binary patterns (LBP) are constructed. For each of the plurality of blocks of the spectrogram, a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks is created. Based on the respective LBP histograms for each of the plurality of blocks of the spectrogram, clusters of LBP histograms are identified, and based upon the clusters, code words for a machine learning model are generated.

[0007] Many embodiments of the present disclosure are directed to apparatus. One apparatus of the present disclosure includes an audio LBP histogram module, a codebook creation module, histogram mapping module, and a support vector machine. The audio LBP histogram module receives an audio signal spectrogram indicative of an audio input. Based on a comparison of adjacent pixels of the spectrogram, the audio LBP histogram module constructs a plurality of local binary patterns (LBP). For each of a plurality of blocks of the spectrogram, the audio LBP histogram creates a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks. The codebook creation module communicatively coupled to the audio LBP histogram module, identifies clusters of LBP histograms, with a k-means algorithm, from the respective LBP histograms. The histogram mapping module communicatively coupled to the codebook creation module and audio LBP histogram module, generates, based upon the clusters, code words which are a centroid of the identified clusters. Based upon correspondence between the LBP histograms and a plurality of code words, the histogram mapping module creates a codebook histogram. Finally, the support vector machine module, communicatively coupled to the histogram mapping module, classifies, using a machine learning model, the codebook histogram to identify context for the received audio signal.

[0008] The above discussion/summary is not intended to describe each embodiment or every implementation of the present disclosure. The figures and detailed description that follow also exemplify various embodiments.

DESCRIPTION OF THE FIGURES

[0009] Various example embodiments may be more completely understood in consideration of the following detailed description in connection with the accompanying drawings, in which:

[0010] FIGS. 1A-B shows functional block diagrams, consistent with various aspects of the present disclosure;

[0011] FIG. 2 shows a system level block diagram, consistent with various aspects of the present disclosure;

[0012] FIG. 3 is a diagram showing the creation of LBPs from an audio signal, consistent with various aspects of the present disclosure; and

[0013] FIG. 4 shows a functional block diagram, consistent with various aspects of the present disclosure.

[0014] While various embodiments discussed herein are amenable to modifications and alternative forms, aspects thereof have been shown by way of example in the drawings and will be described in detail. It should be understood, however, that the intention is not to limit the invention to the particular embodiments described. On the contrary, the intention is to cover all modifications, equivalents, and alternatives falling within the scope of the disclosure including aspects defined in the claims. In addition, the term "example" as used throughout this application is only by way of illustration, and not limitation.

DETAILED DESCRIPTION

[0015] Aspects of the present disclosure generally relate to sound processing. In particular, the present disclosure presents solutions for audio scene recognition via a novel approach utilizing local binary patterns to identify patterns in audio which can be associated with context(s), e.g., different sources, speech, music, background noise and particular events.

[0016] In many applications of computer electronics (e.g., cell phones), customers require high levels of personalization in the data that is presented to the user. Embodiments of the present disclosure are generally directed to enabling such personalization by contextualizing the environment the device is being used in utilizing audio received over a built-in microphone (on-board the electronic device).

[0017] In various embodiments, the device is able to isolate and identify context(s) of the environment and adapt presented information based at least in part on the environment. For example, the device can be configured to determine, based upon classification of sounds in a received audio signal (and other factors such as time of day, location and other device discernable information), that the user was on the street on his way home from work, for example. The device could then display information indicative of current traffic patterns, best routes home, weather conditions, and local restaurants in the area currently offering happy hour specials.

[0018] The context identified by the device may include various distinguishable aspects of the audible environment, such as the hum of fluorescent lights in an office, the exhaust noise of a city bus, or the various ambient noises of a city street (such as car horns), etc. The use of audio scene

recognition over other methodologies provides several benefits including at least the ability to capture temporal events, and the ability to distinguish several co-occurring contexts in the environment. It has been discovered that the ability to correlate a single audio event over a period of time greatly increases the ability to positively identify the audio event (or context). Also, aspects of various embodiments can be particularly useful for limiting additional hardware, as many electronic devices already include microphones and sufficient processing circuitry.

[0019] Aspects of the present disclosure are also directed to reducing processing requirements for the received audio. In many cases, such embodiments are ideally embedded into mobile devices which have a finite amount of electrical power. In further more specific embodiments of the present disclosure, the audio processing methods may further reduce audio processing requirements of the device, thereby increasing battery life if the device is mobile, and improving overall performance of the device as less processing power needs to be allotted to background functions (including audio processing). The present embodiment achieves such a reduction in power usage by minimizing the comparisons required between a histogram associated with a received audio signal and codebook histograms by, for example, clustering the various codebook histograms that are located relative to one another into a single centroid (this also reduces memory usage of the audio processing methodology).

[0020] Aspects of the present disclosure apply image processing techniques on an audio spectrum by transforming the audio into an audio spectrogram (a visual representation of the audio at one or more periods of time) and analyzing a series of spectrograms from the same audio source to classify and identify audio patterns that are indicative of an environmental context that a device receiving the audio is in (e.g., office, bus, street, etc.). The ability to map these audio patterns over time can greatly increase the acoustic recognition performance. In embodiments of the present disclosure, the acoustic recognition method may utilize a pre-defined codebook that can be used to classify and identify audio patterns relative to known environmental contexts. When compared, if the identified audio pattern is within a similarity threshold to one or more of the known audio patterns in the codebook, the identified audio pattern will be associated with the known audio pattern's environmental context. In yet further embodiments, the method may include a learning-phase which allows a user utilizing the acoustic recognition method to identify audio patterns (which the acoustic recognition method may not otherwise be able to currently identify given its existing codebook) and add such contexts to the codebook.

[0021] Various example embodiments of the present disclosure are directed to methods, circuits, and systems that address challenges directed to acoustic context recognition and its implementation. It should also be understood that aspects of the present disclosure may also be utilized in other areas of audio recognition including speech recognition.

[0022] Embodiments of the present disclosure are directed to a learning phase for identifying context(s) of an audio scene. The identified context can later be used in (real-time) applications to identify contexts of a received audio input. According to example embodiments, a processing device (e.g., one or more computer processor units or circuits) receives an audio signal in the form of a spectrogram that is

indicative of an audio input. The processing device can then compare adjacent pixels of the spectrogram to construct a plurality of local binary patterns (LBP). For each of a plurality of sub-blocks of the spectrogram (e.g., sub-blocks that correspond to different audio frequency ranges), a respective LBP histogram is created based on a number of times different LBPs occur. The processing device can then use the respective LBP histograms for each of the plurality of blocks of the spectrogram to identify clusters of LBP histograms. Code words are then created from the clusters (e.g., code words that are based upon an average or centroid of each of the clusters). The processing device code words can then be classified according to a machine learning algorithm or model.

[0023] Various embodiments of the present disclosure are directed to a method for identifying context(s) of an audio scene. According to example embodiments, a codebook that was created based upon clustering of LBP histograms can be used as a feature list. Particular details regarding the creation of such a codebook are described in more detail herein. The use of such a codebook can be particularly useful in that it provides a compact and discriminant set of features from which to classify the audio content. Consistent with embodiments, a processing device can receive an audio signal spectrogram, indicative of an audio input. The processing device can then construct a plurality of LBPs based on a comparison of adjacent pixels of the spectrogram.

[0024] For each of a plurality of sub-blocks of the spectrogram, a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks, is created. The plurality of LBPs are indicative of pixel values, threshold transitions, and respective pixel locations in each of the plurality of blocks of the spectrogram. The respective LBP histogram, for each of the plurality of blocks of the spectrogram, is indicative of an acoustic context of the audio input over a period of time. In more specific embodiments, each respective LBP histogram maps each of the plurality of blocks of the spectrogram into a plurality of uniform configurations and a non-uniform configuration. Based upon correspondence between the LBP histograms and a plurality of code words, a codebook histogram is created. In further embodiments, the creation of the codebook histogram further includes determining a cosine distance between the LBP histograms and the code words using a cosine distance between the LBP histograms and the code words, and assigning the LBP histograms to code words in the codebook histogram based upon the cosine distance.

[0025] Finally, using a machine learning model that was trained using the codebook, the codebook histogram is classified to identify context for the received audio signal. In various embodiments, the machine learning model is within a support vector machine. In yet further more specific embodiments of the present disclosure, the clusters of LBP histograms can be identified using a k-means clustering algorithm. The k-means clustering algorithm is a machine learning technique that partitions n observations into k clusters in which each observation belongs to the cluster with the nearest mean. Given a set of observations (x_1, x_2, \dots, x_n) , where each observation is a d -dimensional real vector, k-means clustering aims to partition the n observations into k ($\leq n$) sets $S = \{S_1, S_2, \dots, S_k\}$ so as to minimize the within-cluster sum of squares (WCSS). The k-means clustering equation is:

$$\operatorname{argmin}_s \sum_{i=1}^k \sum_{x \in S_i} \|x - \mu_i\|^2$$

where μ_i is the mean of points in S_i . The algorithm ultimately finds clusters of a comparable spatial extent. The processing device can then generate code words that are based upon centroids for the clusters.

[0026] Many embodiments of the present disclosure further include interpolating, for each of the LBP histograms, using pixel values outside a respective block for each of the LBP histograms.

[0027] Many embodiments of the present disclosure are directed to apparatus. One apparatus of the present disclosure includes an audio LBP histogram module, a codebook creation module, histogram mapping module, and a support vector machine. The audio LBP histogram module receives an audio signal spectrogram indicative of an audio input. Based on a comparison of adjacent pixels of the spectrogram, the audio LBP histogram module constructs a plurality of local binary patterns (LBP). For each of a plurality of blocks of the spectrogram, the audio LBP histogram creates a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks. The codebook creation module communicatively coupled to the audio LBP histogram module, identifies clusters of LBP histograms, with a k-means algorithm, from the respective LBP histograms. The histogram mapping module communicatively coupled to the codebook creation module and audio LBP histogram module, generates, based upon the clusters, code words which are a centroid of the identified clusters. Based upon correspondence between the LBP histograms and a plurality of code words, the histogram mapping module creates a codebook histogram. Finally, the support vector machine module, communicatively coupled to the histogram mapping module, classifies, using a machine learning model, the codebook histogram to identify context for the received audio signal.

[0028] Turning now to the figures, various embodiments of the present disclosure are presented by way of the illustrations.

[0029] FIG. 1A shows a functional diagram for recognizing an acoustic context of an audio signal, consistent with various aspects of the present disclosure. First, an audio signal spectrogram indicative of an audio input is received **105**. A plurality of local binary patterns (LBP) are then constructed **110** based on a comparison of adjacent pixels of the received spectrogram. For each of a plurality of blocks of the spectrogram, a respective LBP histogram is created **115** based on a number of times different LBPs occur in a corresponding block of the plurality of blocks. From each of the respective LBP histograms, clusters of LBP histograms are identified **120**. For each of the clusters, a code word representing a corresponding cluster is generated **125**. The LBP histograms and the code words representing the clusters are then compared to create a codebook histogram **130**. Finally, using a machine learning model, the codebook histogram is classified **135** to identify context for the received audio signal.

[0030] FIG. 1B shows another functional diagram for recognizing an acoustic context of an audio signal, consistent with various aspects of the present disclosure. First, an

audio signal spectrogram indicative of an audio input is received **150**. A plurality of local binary patterns (LBP) are then constructed **155** based on a comparison of adjacent pixels of the received spectrogram. For each of a plurality of blocks of the spectrogram, a respective LBP histogram is created **160** based on a number of times different LBPs occur in a corresponding block of the plurality of blocks. From each of the respective LBP histograms, clusters of LBP histograms are identified **165**. Based upon the clusters of the LBP histograms, a codebook is generated containing a plurality of code words **170**, then the LBP histograms are mapped to the code words of the codebook to create a codebook histogram **175**. Finally, a machine learning algorithm is trained **180** using the LBP histograms as features therefore.

[0031] FIG. 2 shows a system level block diagram, consistent with various aspects of the present disclosure. The system **200** includes a learning device **201** that creates a (support vector machine, "SVM") model that classifies known contexts by using a codebook of histograms as the features used in the classification. A device **211** receives the (SVM) model from the learning device **201**, and uses the model to classify received sounds for which the context is not known (the real-time usage). In various embodiments of the present disclosure, the model may be transferred to the device **211** utilizing wired or wireless communication means. Furthermore, in some embodiments, the model may be regularly updated to include newly classified audio context.

[0032] In reference to the learning device **201**, a memory module **202** containing audio data with known contexts, provides the audio data to a features extraction module **205**. The features extraction module **205** constructs audio signal spectrograms in a temporal manner over the time-length of the audio data. The spectrograms are then divided into sub-blocks and each sub-block is multiplied by an audio adapted local binary pattern ("LBP"). Histograms of LBPs are extracted for each sub-block, each histogram comprising the audio spectrum of the environment. The features extraction module **205** then utilizes a clustering algorithm to create a codebook of acoustic patterns. These acoustic patterns represent common characteristics (or features) that can be used to identify the context (or noise) that may be relevant for a given application.

[0033] Consistent with embodiments of the present disclosure, the clustering of the LBPs reduces, or compresses, the features for analysis. This can be particularly useful for reducing the use of device processing and memory resources. In particular embodiments, with all the histograms extracted, a k-means clustering algorithm (based on cosine similarity) may be utilized to group similar histogram blocks and to obtain a final codebook with the relevant patterns of the dataset. Centroids of the clusters are written as the elements of the codebook. Supervised training module **210** creates a (SVM) model that classifies the known contexts by using the codebook of histograms as the features used in the classification. The output is an (SVM) model that can be used to classify future sounds for which the context is not known (the real-time usage).

[0034] In reference to the device **211** of the system **200**, the model of the learning device **201** is configured in the features extraction module **206** of the device **211**. The model including a codebook that represents the common characteristic, or features, of the types of context that are desirable

to detect in a given application(s). An audio signal is received by an audio signal receiving module **220** (e.g., microphone, or other audio reception mechanism). The audio signals, in many embodiments, are indicative (at least in part) of an environment in which the audio signal receiving module **220** is located within. For example, if a device such as a cellphone including the various modules of FIG. 2 is at a sporting event, the acoustic context (including, e.g., speech, music, background noise and other particular acoustic events) can be processed according to aspects of the present disclosure to indicate the contextual environment of the cellphone, the sporting event. This can include classifying various audio features according to learned characteristics that are consistent with a sporting event. In such embodiments (and potentially in conjunction with other collected data such as location data, user inputs, etc.), the cellphone can be configured to customize the user interface to display relevant information to the sporting event such as up-to-date player statistics, news, instant replays, and other league related news.

[0035] The audio signal receiving module **220** may conduct a number of ancillary functions to the received audio signal such as noise reduction, amplification, etc. before transmitting the audio signal to the features extraction module **206**. The features extraction module **206** first constructs an audio signal spectrogram indicative of the audio output received by the audio signal receiving module **220**. The features extraction module **206** then develops a plurality of local binary patterns (LBP) based on a comparison of adjacent pixels of the spectrogram. Based on each of a plurality of blocks of the spectrogram, a respective LBP histogram is created that is based on the number of times different LBPs occur in a corresponding block of the plurality of blocks. Based upon correspondence between the LBP histograms and a plurality of code words, a codebook histogram is created. Finally, the recognition SVM **225** classifies, using a machine learning model, the codebook histogram to identify context for the received audio signal. The recognition result **230**, provided as an output of recognition SVM **225**, can indicate the context(s) for the received audio. It is to be understood that other classifiers used in supervised learning may also be utilized in the present disclosure, besides a support vector machine, e.g., artificial neural networks, and Gaussian mixture models.

[0036] FIG. 3 shows a diagram showing the creation of LBPs from an audio signal, consistent with various aspects of the present disclosure. In various embodiments, a processing device **300** can include a features extraction module that can be used to generate LBPs from a spectrogram **330** of an audio signal. A possible format for the spectrogram **330** is a graph where one axis (e.g., the horizontal axis) represents time and another axis (e.g., the vertical axis) represents frequency. The amplitude or volume can be indicated by the intensity or color of the image pixels.

[0037] The features extraction module can break the spectrogram **330** into a number of sub-blocks, including block **331**. The spectrogram **330**, in certain embodiments, can be a linear-spectrum spectrogram. Due to the local nature of the sub-blocks, each sub-block can represent a particular time and frequency range of the audio signal. It has been discovered that linear representations of a spectrogram contain less noise and can be primarily dominated by sparse high-energy elements (increasing the likelihood of pattern recognition). As shown in **332**, for each pixel of the image, a block is built

around the central pixel. As shown in **333**, the value of each pixel can be compared with the surrounding pixels (using thresholding) if the values are lower, they are mapped as zeroes; if greater, as ones. In specific embodiments, the value of the surrounding pixels can be compared with a mean over the entire block **331**, although other variations for determining a comparison threshold are possible. In such embodiments, this mean approach can make the algorithm more robust to noise and to large variations between pixels.

[0038] Starting from the top-left corner of block **332**, the binary values are transformed (see, e.g., Binary: 11010011). In certain embodiments the binary result of the block **332** can be modeled as indicative of one of the 59 possible bit combination circles depicted in **340** by interpolating the values. In particular, the first 58 bit combinations represent uniform bit combinations in which the number of binary transitions between ones and zeroes is no greater than 2. A non-uniform LBP can be mapped to a 59th bit combination. Accordingly, the dimensions of the resulting histogram **350** can be limited to one of the 59 possible patterns in the algorithm (each pattern taking into account the possible orientations of the resulting block **333**). Consistent with various embodiments, other histograms are also possible including more exhaustive lists of bit combinations or further reduced bit combinations. For example, one or more non-uniform bit combinations may be particularly discriminatory in certain applications. The histogram dimensions could therefore be expanded to include such bit combinations. Alternatively, or in addition, one or more uniform bit combinations may be removed for certain applications.

[0039] Each time a pattern is found in the block (one of the 59 possible patterns), that pattern is accumulated in the resulting histogram **350**. The resulting histogram **350** takes into account the occurrences of each i-th pattern and directly represents the pattern distributions of the image.

[0040] FIG. 4 shows a functional block diagram for recognizing an acoustic context of an audio signal, consistent with various aspects of the present disclosure. The solution of FIG. 4 exploits an LBP, an image processing technique that utilizes an audio signal spectrogram representation, a histogram, to recognize an acoustic context. The LBP can be used as an efficient texture operator which labels the pixels of a spectrogram image by thresholding the neighboring pixels and computing the result as a binary number (as discussed in more detail herein).

[0041] The LBP is particularly well-suited to various audio applications as its algorithm is robust even against huge fluctuations between neighboring pixels in a spectrogram (which are common in audio spectrograms). The LBP also benefits from the ability to build more discriminative features, which utilize the time-frequency information of the spectrum. It has been discovered that capturing and correlating context events over a period of time can significantly increase the recognition performance of the context event.

[0042] Exemplary circuitry to recognize acoustic contexts using an LBP-based algorithm are described in detail in reference to FIG. 4. In accordance with embodiments of the present disclosure, the circuitry is partitioned into five blocks/modules: audio LBP histograms module **1**, codebook creation module **2**, histograms mapping module **3**, SVM training module **4**, and SVM testing module **5**.

[0043] The audio LBP histograms module **1** receives an audio signal **401** and converts the audio signal into a spectrogram **402**. The spectrogram is then divided into

blocks **403** (e.g., blocks that correspond to different audio frequency ranges) and utilizing an LBP algorithm finds recurrent patterns from the input image (spectrogram) by comparing the value of each pixel with its respective neighbors. A LBP histogram **404** is then built for each of the blocks **403** of the spectrogram **402**, counting how many times a pattern is found in the entire image. The input image can be a spectrogram **402** of an audio signal **401** and it has to be modified to better fit with the LBP algorithm.

[0044] During supervised training of the system using known audio sequences, the codebook creation module can be used to identify features that can then be used to train a machine learning module, such as an SVM. Once the LBP histograms from the entire dataset have been extracted, the codebook creation module **2** can employ a clustering technique **405** to group together frequently used histograms. The cluster values can then be used to form a codebook **406**. The histograms mapping module **3** then exploits the cosine similarity **407** to find which codebook histograms **408** in the codebook **406** are similar. The LBP histograms **404** can then be mapped onto one of the code words from the codebook. After these steps, the features of the codebook histograms **408** are less redundant and with lower dimensions than the LBP histograms **404** alone. This can be particularly useful for embedded devices, which may have memory and computational constraints.

[0045] The codebook histograms **408** from the histograms mapping module **3** become the input features of the SVM training module **4** and SVM testing module **5**. In various embodiments of the present disclosure, the SVM training module **4** will be trained **409** to classify either on the 59-dimensions histograms for the three sub-bands or on the features mapped into one of the possible pattern(s) in the dictionary. As a result, an SVM model **410** is created, which can be used by the SVM testing module **5** to predict **411** and output **412** a context indicative of a received audio signal. In many embodiments the SVM modules **4** and **5** project the data into a higher dimensional space, in this new space it is possible to apply a linear separation using the hyperplane with the maximal margin and the minimum error of classification. In addition, a denser representation is able to maximize the distance between the different contexts and improve the discriminative power of SVM.

[0046] In certain experimental embodiments of the present disclosure, an LBP algorithm is disclosed that compares the surrounding pixels with the mean over the entire block:

$$LBP_{p,R} = \sum_{i=0}^{p-1} f(g_i - \mu) 2^i,$$

$$f(x) = \begin{cases} 1, & x \geq \mu \\ 0, & x < \mu \end{cases}$$

where g_i is the i-th neighbour's value, μ is the mean over the block, P is the number of involved pixels. R is the radius of the neighbourhood: the coordinates of g_i are $R \cos(2\pi i/P)$, $\sin(2\pi i/P)$. The pixel values that are not in the block can be estimated by bi-linear interpolation. A pixel can be then interpolated using both x and y coordinates in the 3*3 block combined with interpolation weights w_i :

$$z = w_0 + w_1x + w_2y + a_3xy$$

[0047] In various example embodiments, the pixel values in the spectrogram are affected by huge fluctuations which can detriment LBP representation (resulting in a possibly limitless number of potential dimensions of a histogram). To reduce the potential dimensions of a histogram, consider the transition between 0's and 1's in the LBP code: if the number of transitions is less or equal to 2, the LBP binary string is considered uniform and it will be mapped into one of 58 configurations, like in FIG. 2 (the 59th is therefore for the non-uniform ones). Uniform patterns ensure the presence of a relevant texture element, such as an edge, corner or uniform zone and it will discard the non-uniform ones (to the 59th configuration), which are more likely to have been affected by noise. In experimental testing, this histogram reduction has been demonstrated to perform well.

[0048] Referring to spectrogram analysis in various embodiments of the present disclosure, pixels in the spectrogram express energy information in specific time-frequency coordinates. Many embodiments of the present disclosure utilize linear spectrum spectrographs to log-spectrum spectrographs, which are less dominated by noise. The linear spectrograms are primarily formed by sparse high-energy elements, while the remaining elements do not disturb pattern recognition. In such embodiments, frequency-bands are separated in three different bands, less than 900 Hz, from 900 to 2 kHz and from 2 kHz until 8 kHz, to represent the entire information contained in the signal, sampled at 16 kHz. The particular frequency ranges can be adjusted based upon the particular sampling rate and other considerations, such as the types of context being analyzed. The LBP algorithm uses these three sub-bands to perform the acoustic pattern recognition, and to extract a histogram for each of the corresponding blocks of the spectrogram (like in FIG. 2).

[0049] In specific embodiments of the present disclosure, the algorithm used for creating the codes in the codebook includes a further step before the classification to make the features more compact. This further step discovers the most representative patterns in the dataset and extracts them automatically, using an unsupervised classification. The creation of a codebook of the most representative histograms is an essential part of the algorithm, as it allows for the classification of unknown acoustic contexts. A k-means clustering algorithm (or other clustering algorithm) may be used to group the similar blocks and obtain a final codebook with the most relevant patterns of the dataset. Cosine distance may be used as a good metric for many feature descriptors, especially histogram features. Finally, the centroids of the clusters become the elements of the codebook, as described in more detail above.

[0050] Various blocks, modules or other circuits may be implemented to carry out one or more of the operations and activities described herein and/or shown in the figures. In these contexts, a "block" (also sometimes "logic circuitry" or "module") is a circuit that carries out one or more of these or related operations/activities (e.g., first module, second module, and memory module). For example, in certain of the above-discussed embodiments, one or more modules are discrete logic circuits or programmable logic circuits configured and arranged for implementing these operations/activities, as shown in FIG. 1. In certain embodiments, such a programmable circuit is one or more computer circuits programmed to execute a set (or sets) of instructions (and/or configuration data). The instructions (and/or configuration

data can be in the form of firmware or software stored in and accessible from a memory (circuit). As an example, first and second modules include a combination of a CPU hardware-based circuit and a set of instructions in the form of firmware, where the first module includes a first CPU hardware circuit with one set of instructions and the second module includes a second CPU hardware circuit with another set of instructions.

[0051] Certain embodiments are directed to a computer program product (e.g., nonvolatile memory device), which includes a machine or computer-readable medium having stored thereon instructions which may be executed by a computer (or other electronic device) to perform these operations/activities.

[0052] Based upon the above discussion and illustrations, those skilled in the art will readily recognize that various modifications and changes may be made to the various embodiments without strictly following the exemplary embodiments and applications illustrated and described herein. For example, acoustic context recognition, as disclosed herein, could be utilized in a speech recognition environment, and/or other computer-based audio recognition applications that would benefit from the various advantages of the present disclosure. Such modifications do not depart from the true spirit and scope of various aspects of the invention, including aspects set forth in the claims.

1. A method comprising:

- receiving an audio signal spectrogram, having adjacent pixels, indicative of an audio input;
- constructing a plurality of local binary patterns (LBP) based on a comparison of the adjacent pixels of the spectrogram;
- creating, for each of a plurality of blocks of the spectrogram, a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks;
- identifying clusters of LBP histograms from the respective LBP histograms;
- generating, for each of the clusters, a code word representing a corresponding cluster;
- creating a codebook histogram based upon correspondence between the LBP histograms and the code words representing the clusters; and
- classifying, using a machine learning model, the codebook histogram to identify context for the received audio signal spectrogram.

2. The method of claim 1, wherein the respective LBP histogram, for each of the plurality of blocks of the spectrogram, is indicative of an acoustic context of the audio input over a period of time.

3. The method of claim 1, wherein the audio signal spectrogram is a linear spectrum representation of the audio input.

4. The method of claim 1, wherein the plurality of LBPs are indicative of pixel values, threshold transitions, and respective pixel locations in each of the plurality of blocks of the spectrogram.

5. The method of claim 1, wherein creation of the codebook histogram further includes:

- determining a cosine distance between the LBP histograms and the code words using the cosine distance between the LBP histograms and the code words, and
- assigning the LBP histograms to code words in the codebook histogram based upon the cosine distance.

6. The method of claim 1, further comprising the step of interpolating, for each of the LBP histograms, by using pixel values outside a respective block for each of the LBP histograms.

7. The method of claim 1, wherein each respective LBP histogram is a histogram of a relationship between each of the plurality of blocks of the spectrogram as defined by a plurality of uniform configurations and a non-uniform configuration.

8. The method of claim 1, wherein the plurality of blocks of the spectrogram are defined by blocks within 0 Hz-900 Hz frequency-band, blocks within 900 Hz-2000 Hz frequency-band, and blocks within 2000 Hz-8000 Hz frequency-band.

9. The method of claim 1, wherein the LBP is defined by an algorithm:

$$LBP_{P,R} = \sum_{i=0}^{P-1} f(g_i - \mu)2^i,$$

$$f(x) = \begin{cases} 1, & x \geq \mu \\ 0, & x < \mu \end{cases},$$

wherein g_i is a neighboring pixel's value, μ is a mean over a particular block, P is a number of pixels in the particular block, R is a radius of a neighborhood and wherein coordinates of g_i are $R \cos(2\pi i/P)$, $\sin(2\pi i/P)$.

10. A method comprising the steps of:
 receiving an audio signal spectrogram indicative of an audio input;
 constructing a plurality of local binary patterns (LBP) based on a comparison of adjacent pixels of the spectrogram;
 creating, for each of a plurality of blocks of the spectrogram, a respective LBP histogram that is based on a number of times different LBPs occur in a corresponding block of the plurality of blocks;

identifying clusters of LBP histograms from the respective LBP histograms for each of the plurality of blocks of the spectrogram;

generating, based upon the clusters, a codebook containing a plurality of code words;

mapping the LBP histograms to the code words of the codebook to create a codebook histogram; and
 training a machine learning algorithm using the LBP histograms as features therefore.

11. The method of claim 10, further comprising interpolating, for each of the LBP histograms, using pixel values outside a respective block for each of the LBP histograms.

12. The method of claim 10, wherein each respective LBP histogram maps each of the plurality of blocks of the spectrogram into a plurality of uniform configurations and a non-uniform configuration.

13. The method of claim 10, wherein the plurality of blocks of the spectrogram are defined by blocks within 0 Hz-900 Hz frequency-band, blocks within 900 Hz-2000 Hz frequency-band, and blocks within 2000 Hz-8000 Hz frequency-band.

14. The method of claim 10, wherein the machine learning algorithm is within a support vector machine.

15. The method of claim 10, wherein the LBP is defined by an algorithm:

$$LBP_{P,R} = \sum_{i=0}^{P-1} f(g_i - \mu)2^i,$$

$$f(x) = \begin{cases} 1, & x \geq \mu \\ 0, & x < \mu \end{cases},$$

where g_i is a neighboring pixel's value, μ is a mean over a particular block, P is a number of pixels in the particular block, R is a radius of a neighborhood and wherein the coordinates of g_i are $R \cos(2\pi i/P)$, $\sin(2\pi i/P)$.

* * * * *