

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property  
Organization

International Bureau

(43) International Publication Date  
20 August 2020 (20.08.2020)



(10) International Publication Number  
**WO 2020/168352 A1**

(51) International Patent Classification:

G06K 9/62 (2006.01) G06N 20/00 (2019.01)

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

(21) International Application Number:

PCT/US2020/018697

Published:

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

(22) International Filing Date:

18 February 2020 (18.02.2020)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

16/277,993 15 February 2019 (15.02.2019) US

(71) Applicant: **SHOTSPOTTER, INC.** [US/US]; 7979 Gateway Boulevard, Suite 210, Newark, CA 94560 (US).

(72) Inventors: **CALHOUN, Robert B.**; c/o ShotSpotter, Inc., 7979 Gateway Boulevard, Suite 210, Newark, CA 94560 (US). **LAMKIN, Scott**; c/o ShotSpotter, Inc., 7979 Gateway Boulevard, Suite 210, Newark, CA 94560 (US). **RODGERS, David**; c/o ShotSpotter, Inc., 7979 Gateway Boulevard, Suite 210, Newark, CA 94560 (US).

(74) Agent: **SCHWAAB, Andrew**; c/o Greenberg Traurig, LLP, 77 West Wacker Drive, Suite 3100, Intellectual Property Department, Chicago, IL 60601 (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, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, WS, ZA, ZM, ZW.

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

(54) Title: SYSTEMS AND METHODS INVOLVING CREATION AND/OR UTILIZATION OF IMAGE MOSAICS

(57) Abstract: Systems and methods that yield highly-accurate classification of acoustic and other non-image events are disclosed, such as those involving pre-processing data from one or more transducers and/or generating a visual representation of the source as well as associated features and processing. In some implementations, such pre-processing may be utilized when 1) all impulsive acoustic events have many features in common due to their point source origin and/or impulsive nature, and/or 2) the error rates that are considered acceptable in general purpose image classification are much higher than the acceptable levels in automatic impulsive incident classification. Further, according to some aspects, data may be pre-processed in various ways, such as to remove extraneous or irrelevant details and/or perform any required rotation, alignment, scaling, etc. tasks, such that these tasks do not need to be otherwise "learned" via the neural network.



WO 2020/168352 A1

## SYSTEMS AND METHODS INVOLVING CREATION AND/OR UTILIZATION OF IMAGE MOSAICS

### CROSS-REFERENCE TO RELATED APPLICATION(S)

**[0001]** This is an International Application claiming benefit and priority of U.S. application No. 16/277,993, filed February 15, 2019, now U.S. patent No. 10,424,048, which are incorporated herein by reference in entirety.

### BACKGROUND

#### *Field*

**[0002]** This specification relates to the processing of acoustic and non-acoustic data such as that obtained from one or more acoustic sensor devices and/or other sources and identifying/classifying the nature of the event or source responsible for the data/signal.

#### *Description of the Related Technology*

**[0003]** The use of deep convolutional neural networks in practical applications is becoming synonymous with our daily lives. From self-driving cars to fingerprint recognition (Google's patent US5825907A). These applications rely on real world input and that input is only somewhat pre-processed or uses neural networks to do the bulk of the preprocessing (Christian Szegedy & Google's patent US20160063359A1). Other patents are related primarily with the recognition of speech patterns (Google's patent US20160322055A1 or US20160284347A1). These are insufficient for addressing audio signals with minimal information, similar to using the discrepancies between glottal stops made by different people and classifying the size of the glottis based on the sound. There is information loss involved in the process as well as a limiting of scope when the derivation of an answer is based solely on listening. Some representative architectures, in this regard, include the ResNet and Inception convolutional neural network architectures.

**[0004]** Illegal discharge of firearms or other explosive devices is an ongoing problem in our society. Specific examples include: firearms discharged during the commission of a felony; discharge of a firearm in an unsafe manner that can result in the accidental injury or

death of a person, such as celebratory gunfire; poaching of endangered species or other illegal hunting; blast fishing; and acts of war or terrorism. A system that can detect and accurately locate such activities in near real-time aids law enforcement and other first responders by reducing the time required to arrive at the location of the discharge. A rapid response to such illegal activities assists law enforcement in apprehending the responsible parties and improves community safety by discouraging others from engaging in the illegal activity.

**[0005]** The location of acoustic events can be computed via a variety of methods. Such methods include multilateration (location computation using time difference of arrival); triangulation (location using the intersection of two or more direction of arrival vectors); range estimation using the difference between the optical and acoustic arrival times of a signal; range estimation using signal amplitude measurements along with an estimate of the emitter strength; and combinations of these techniques. Discharge of a firearm and detonation of an explosive device both generate impulsive signals of high amplitude that can be successfully located using these techniques.

**[0006]** A challenge to practical gunshot or explosive detection systems is the presence in the environment of other sources of loud, impulsive signals of lesser or negligible law enforcement concern, such as discharge of fireworks, vehicle backfires, pile drivers or other construction noise, or even the emptying of large refuse-storage containers. A system that seeks to dispatch law enforcement personnel to the scene within a few seconds must necessarily be automated, and in order for the automated system to be successful the false positive rate for automatically dispatched incidents (defined as one minus the ratio of events of law enforcement concern divided by the total number of dispatched events) must be kept to a low value.

**[0007]** Thus, the goal of automated classification of impulsive events is to correctly classify events of unknown origin into groups or classes, such as “gunfire”, “fireworks”, “backfire” or “construction noise”. There is a further goal of identifying subclasses of these classes, such as “discharge of an AK-47” or “discharge of a 45-caliber revolver” that can be used to discriminate between friendly and hostile weapon fire events, or to inform the appropriate magnitude of law enforcement response. In general, numerous techniques exist to classify the source of an acoustic event in an automated manner.

**[0008]** However, these techniques typically do not provide sufficient classification accuracy in complex environments. Impulsive acoustic events that are fundamentally explosive in nature have low information content in comparison with speech, music, bird songs and other structured audio signals. The major components of impulsive acoustic events can be modelled with reasonable accuracy using a single-parameter model, as known in the art since at least 2002. The waveform detected from a single impulsive event is more characteristic of the filter applied by the environment (trees, buildings, etc.) than the source itself. For this reason, secondary characteristics such as the degree of propagation anisotropy, spacing between shots, presence or absence of echoes due to the height above ground and roofline level of the discharge provide valuable additional data that can be used to improve the accuracy of machine classification results. These data are not, in general, independent; for example, discharge of a firearm generates both an impulsive acoustic signal and an impulsive optical signal, the latter being strongest in the medium-wave (2-5  $\mu\text{m}$ ) infrared region of the light spectrum, and the fusion of results from acoustic and optical sensors can be expected to be more accurate than either acoustic or optical sensors alone. For this reason, it is desirable to have a general-purpose method for incorporating data collected from one or more signal transducers into a single large classification vector than can be used as input into a general-purpose machine learning / machine classification system.

**[0009]** One such machine classification system that has made rapid improvements in classification accuracy in recent years are “deep learning” neural networks. These neural networks make use of a large number of intermediate layers levels and may utilize newer training techniques (such as adverse training REF and obfuscation REF) with large datasets to yield impressive accuracy at common tasks such as optical character recognition or automated image classification. Effective Deep Learning neural networks require an intimidating numbers of internal connections. Further, training such networks in practical amounts of time typically requires use of complexity-reducing methods such as a convolutional in combination with parallel training on large numbers of CPUs, GPUs, or even custom ASICs optimized for the training task. These techniques are especially well-developed for the task of image recognition, since images of even modest size imply large feature spaces. For example, a 256 x 256 pixel RGB color image can be considered as a classification vector of length 196,608 where each entry can take one of 256 values for each

color. Rarely, however, are all parts of the image of equal interest. For example, an image containing a face might also contain uninteresting information such as the texture of the wall behind the person. While a fully-connected neural network could in theory learn to recognize a feature such as a face without any down-sampling or simplification steps, a more practical approach is to use a system where the upper levels of the neural network scale, translate, and, if necessary, rotate the salient details of the image so that scaled, rotated and normalized features are passed to the lower levels of the neural network. Such networks are known as Convolutional Neural Networks (CNNs) and were first described by Yann LeCun *et al.* in 1999, *Object Recognition with Gradient-Based Learning*. In a convolution step, a 2D kernel (often itself a product of machine learning) is applied to the 2D input at a particular layer and only the output of the convolution is passed to the next layer. Deep learning networks targeted at image classification are now available as both stand-alone packages and as software-as-a-service offerings.

#### SUMMARY OF CERTAIN ASPECTS

[0010] Because these sophisticated general-purpose image classifiers are widely available, development of systems and methods by which data that is not intrinsically visual in nature can be converted to a visual representation for use in a general-purpose image classifier, as set forth herein, has proven to possess great novelty and utility. The present inventions disclose such systems and methods that yield highly-accurate classification of acoustic events by pre-processing data, such as data from one or more transducers and/or data associated with a subject event or source, and generating a visual representation of the source. According to implementations herein, such pre-processing steps may be utilized in that 1) all impulsive acoustic events have many features in common due to their point source origin and impulsive nature, and 2) the error rates that are considered acceptable in general purpose image classification (e.g., top 5 error rate REF, etc.) are much higher than the acceptable levels in automatic impulsive incident classification. Given such factors/reasons, it is typically desirable to pre-process the data in various ways, such as to remove extraneous or irrelevant details and perform any required rotation, alignment, and scaling tasks, *inter alia*, by using methods based on, derived from and/or bearing some relation to existing

techniques, such that these tasks do not need to be “learned” in a less direct and more expensive manner in the neural network.

### **BRIEF DESCRIPTION OF THE DRAWINGS**

**[0011]** FIG. 1 is a flow diagram illustrating the process of classifying an event caught across multiple transducers, consistent with certain aspects related to the innovations herein.

**[0012]** FIG. 2 is an example of one of the possible implementations of an image mosaic utilizing the time and frequency domains, consistent with certain aspects related to the innovations herein.

**[0013]** FIG. 3 is a flow diagram illustrating the process of creating a usable dataset and training on a convolutional neural network, consistent with certain aspects related to the innovations herein.

**[0014]** FIG. 4 is an illustration of a generic mosaic suitable for use in classification of acoustic events, consistent with certain aspects related to the innovations herein.

**[0015]** FIGs. 5A-5B are illustrations of exemplary mosaics suitable for use in classification of acoustic events, consistent with certain aspects related to the innovations herein.

**[0016]** FIG. 6 is an illustration of another exemplary mosaic suitable for use in classification of acoustic events, consistent with certain aspects related to the innovations herein.

**[0017]** FIG. 7 is an illustration of an exemplary system adapted to perform various features and/or functionality of the disclosed technology, consistent with certain aspects related to the innovations herein.

### **DETAILED DESCRIPTION OF CERTAIN ILLUSTRATIVE EMBODIMENTS**

**[0018]** Aspects of the disclosed technology relate to systems and methods for pre-processing data received from one or more transducers in ways conducive to training a Deep Learning convolutional neural network classifier optimized for image classification.

**[0019]** With regard to the source location problem, it is desirable to have a plurality of devices at distinct, known locations such that time difference of arrival, angle of

arrival, and/or acoustic-optical travel time distance can be used to accurately locate the source of the acoustic event and thus estimate the range from the event to each reporting sensor device.

**[0020]** FIG. 1 is a flow diagram illustrating the process of classifying an event caught across multiple transducers, according to various aspects consistent with the innovations herein. According to a first part or phase of the implementation(s) illustrated in FIG. 1, at 110, signal processing is performed on said remote sensor device to look for acoustic events of interest. In some aspects of the present inventions, when such acoustic events of interest are identified, further signal processing may be performed to reduce the dimensionality of the input transducers signals so that the resulting feature vector may be easily transferred over a low-bandwidth communication channel, such as a wireless network.

**[0021]** In a second part or phase of the implementation(s) illustrated in FIG. 1, at 120, these low dimensional data from one or more remote sensing devices may be utilized to locate the acoustic event in space and time using source location techniques known to those skilled in the art.

**[0022]** In a third part or phase of the implementation(s) illustrated in FIG. 1, at 130, a coarse identification of the location of the acoustic event may be determined using the low-dimensional data in conjunction with a machine classifier system that returns a numerical estimate that the event is of a type of interest to the system. Many such classifier systems are known to those skilled in the art and include but are not limited to naive Bayesian classifiers, neural network classifiers, convolutional neural network classifiers, support vector machines, regression classifiers, or k-nearest-neighbor classifiers.

**[0023]** In a fourth part or phase of the implementation(s) illustrated in FIG. 1, at 130, the quantitative output value from the third part of the invention, such as a confidence, likelihood, or probability value that the event in question is one of interest (such as a firearm discharge or explosion) is analyzed to determine a subsequent/appropriate course of action. According to some aspects, for example, the quantitative output value may be compared with a fixed or dynamically determined threshold value. Those events for which the output value falls below the threshold are discarded, while additional processing is performed on those events for which the output value is greater or equal to the threshold.

**[0024]** In the fifth part or phase of the implementation(s) illustrated in FIG. 1, at 140, more detailed signal data is obtained/generated/made and processed for those events that meet or exceed the threshold. For example, several seconds of waveform data collected at a sampling rate from as 4 to 48 kHz and a dynamic range of 8 bits to 32 bits from each transducer involved in the detection could either be sent directly to the processing unit performing the machine classification step, or retrieved from the remote device on demand. Further, here, various lossy and lossless compression algorithms can be used to reduce the bandwidth required to return these signal waveforms while preserving their essential characteristics.

**[0025]** In the sixth part or phase of the implementation(s) illustrated in FIG. 1, at 140 and/or 160, the signals are down-sampled, transformed, translated, scaled, and/or converted to bitmaps such that one or more consistent visual representations of each transducer signal are generated.

**[0026]** In the seventh part or phase of the implementation(s) illustrated in FIG. 1, at 150, non-acoustic signal data, such as a computed distance from the acoustic source to the receiver, the history of recent acoustic events in the same area as the event just detected, local elevation information, granular weather information, etc., may be used in combination with transducer signal data (such as peak amplitude, or frequency distribution) and converted into visual representation such as plots, graphs, charts, or diagrams that can assist with correct identification.

**[0027]** In the eighth part or phase of the implementation(s) illustrated in FIG. 1, at 160, these multiple visual representations of data derived from one or more transducers are combined into a mosaic or composite image in a pre-determined configuration. Because the mosaic image comprises data from the same acoustic event, as detected on one or more transducers, and as processed using one or more signal processing techniques, the image generated is a composite image comprising data obtained from of a plurality of remote devices, each of which comprises one or more transducers, converted into visual representations using a plurality of visual representations, such as waveforms, spectrograms, power spectra, plots, graphs, charts and similar methods for visualizing quantitative information.

**[0028]** In the ninth part or phase of the implementation(s) illustrated in FIG. 1, the mosaic tile image generated using the processing steps above is passed, at 165, to at least one convolutional neural network and/or other automated image classifier. Further, in some aspects, such automated image classifier may have previously been trained on images prepared in this way. The automated image classifier may then synthesize the data from these several transducers and generate more accurate results than if each transducer and processing step were to be considered independently. As such, at a high level, innovations herein may be seen as disclosing techniques of “sensor fusion” in which the combined output of multiple transducers located on one or more remote receivers yields superior results in comparison with analysis of waveform or waveform image from any single transducer.

**[0029]** Systems and methods herein allow input from multiple transducers located at the same or different positions to be incorporated into a single image in ways that yield increased performance over a single image. The data is pre-processed, such as to align and scale image, so as reduce training time and network complexity since rotational invariance is not required.

**[0030]** In the tenth part or phase of the implementation(s) illustrated in FIG. 1, at 170, multiple machine classification systems may be executed in parallel, providing redundancy in the case of failure at any stage and to provide a more robust result in the case that they all return a response. This part of the process is called the ensemble and entails using multiple outputs of classifiers. These outputs are then appropriately converted to representative confidence values, with order maintained, and are either combined appropriately or referenced at this point.

**[0031]** FIG. 2 is an example of one of the possible implementations of an image mosaic utilizing the time and frequency domains, consistent with certain aspects related to the innovations herein. FIG. 2 demonstrates one illustrative combination of mosaic elements, with 210 being the wavelet spectrogram, 220 being the waveform, 230 being the power spectra, and 240 being the spectrogram. The alignment of 210 and 220 with respect to the time domain and the alignment of the power spectra and the spectrogram with respect to the frequency domain are examples of mosaic organization/layouts configured to advantageously utilize the convolutional aspect of the neural network.

**[0032]** Prior to utilizing a neural network, the neural network must be trained on a substantial set of data. Such a process of “training a neural network” involving acoustic data is illustrated in FIG. 3.

**[0033]** In FIG. 3, at 310, the database is used to retrieve data pertinent to a unique incident. This data will eventually become a member of the dataset after some manipulation.

**[0034]** Next, at 320, the acoustic and non-acoustic data are retrieved from the database, with labels assigned by our Service Operations Center. This data would need to be generated at the Location, Initial Classifier, and Additional Signal Data stages in FIG. 1 at runtime, though implementations herein retain the information after runtime incidents.

**[0035]** At 330, this data is converted into image representations as described above in connection with FIGs. 1-2, to compile the mosaic. The mosaics/data are then partitioned into two datasets for one for training and for validation testing, at 340.

**[0036]** The training dataset may then be used, at 350, to train the model with a convolutional neural network blueprint. Step 350 may be performed via established practices for training any neural network known to those of skill in the art. Model artifacts, at 360, are the checkpoints for the model at different moments in the training, we use the artifacts that produce the best results at step 380. Here, for example, empirical analyses involving the models and artifacts used to validate accuracy with respect to the audio data processing disclosed herein may be utilized to identify such best results, e.g., to aid in determining which of the various audio features in the underlying mosaics contribute to accurate training/matching. After hosting multiple sets of artifacts or checkpoints at a suitable host, e.g., at Amazon Web Services (AWS), etc., the accuracy and confusion matrix may be validated against the testing dataset, at 370.

**[0037]** FIG. 4 is an illustration of a generic mosaic 400 suitable for use in classification of acoustic events, consistent with certain aspects related to the innovations herein. Representative graphical elements of a generic mosaic are depicted in FIG. 4, including components of one or more domain-related fields 410, one or more fields or rows containing additional images 420, and one or more fields or rows containing visual representations 430 of metrics, features or other graphical indicia regarding the subject audio events. Referring to the domain-related field(s), the mosaic, or portions thereof, may be characterized by a certain domain, such as timeframe, frequency, area, distance-frequency,

distance-timeframe, volume, power, energy, etc. For example, the mosaic may be constructed such that various mosaic elements are mapped against one or more of such common domain(s). These domains include any representation of a higher domain projected down to a lower dimension in addition to any two or lower dimensional domains, which can be directly represented. For example, the x-axis may be time or frequency domain corresponding to the subject events, and the mosaic may contain various subcomponents mapped against such shared domain. Referring to FIG. 4, here, the generic mosaic is shown having a plurality of subcomponents (412A, 412B .. 412n) that are mapped/matched against the respective domain. Further, as explained in more detail below, such subcomponents 412 may include one or more graphs, spectrograms, outputs, standardized outputs, and/or other graphical depictions of the acoustic event mapped against the selected domain. As a result of this graphical comparison over a domain, such mosaic(s) may have the effect of creating interactions between any combination of the graphs in order to create greater possibilities for the neural network to identify and/or latch onto.

**[0038]** While a specific image and component order (“composition order”) is depicted and described in connection with FIG. 4, this particular order shown is not necessarily needed to take advantage of the effects of the convolution neural network. Other composition orders can be utilized to emphasize other aspects and just need to be used consistently between training and practice. Such differing order or arrangements each create a unique perspective in the data when convoluted with its neighbors.

**[0039]** In the generic mosaic 400 of FIG. 4, a field or row containing additional images 420 is also shown. While in FIG. 4 this set of additional images 420 is depicted as a row of images 422A, 422B, 422C ... 422n below the domain-related field 410, such sets of additional images may be positioned or arranged in differing graphical relation to other components of a mosaic. These additional images may include though are not limited to graphs, waveforms, diagrams, graphical indicia of potentially related incidents or events, features with respect to a parameter (e.g., distance, etc.) of an event, other images, and/or portions or snippets thereof.

**[0040]** Further, the generic mosaic 400 of FIG. 4 may also include one or more rows or fields containing additional visual representations 430, such as visual representations of metrics, spectra, other visual indicia, and/or versions of components shown

in the other fields of the mosaic (e.g., fields 410 and 420). Here, *inter alia*, one objective is to get certain new features for the neural network to latch onto after the convolutions have blurred and distorted the image, so a repetition of cells is encouraged, e.g., in row 430. While in FIG. 4 these additional visual representations 432 are depicted as a row of subcomponents or images 432A, 432B, 432C, 432D, 432E ... 432n below the row containing additional images 420, such additional visual representations 432 may be positioned or arranged in differing graphical relation to other components of a mosaic.

**[0041]** FIGs. 5A-5B and 6 depict several representative examples of illustrative mosaics that may be utilized to graphically characterize an audio impulse or signal. While these examples represent some implementations of the presently-disclosed mosaics, it should be noted that many other possible combinations or permutations of the mosaic innovation(s) are set forth and covered via this disclosure and these examples are not limiting to other such mosaic innovations that are described and claimed herein.

**[0042]** FIGs. 5A-5B are illustrations of exemplary mosaics suitable for use in classification of acoustic events, consistent with certain aspects related to the innovations herein. Illustrative graphical elements of the exemplary mosaic are depicted in FIGs. 5A-5B, and are described next, in order, from top to bottom, then left to right. The mosaic, or portions thereof, may be characterized by a certain domain, such as timeframe, frequency, area, etc. For example, the mosaic may be constructed such that various mosaic elements are mapped in or against one or more of such common domain(s). For example, the x-axis may be the time domain corresponding to the subject events, and the mosaic may contain various subcomponents mapped against such shared time domain. Referring to FIGs. 5A-5B, here, the illustrative mosaics are shown having 5 exemplary subcomponents (510, 520, 530, 540 and 550) that are mapped/matched against the respective time domain. Specifically, a first element of the mosaic, shown at the top of FIG. 5A, may comprise a graphical depiction of the audio signal from the sensor adjusted through a frequency domain filter 510. Below this, at 520, the mosaic may include the wavelet spectrogram of the signal. Next, at 530, a standardized output, such as an output that uses the input from multiple sensors to create a simple representation in the time domain, may be included. Further, at 540, another wavelet spectrogram of the noise, rather than the signal, may be shown as mapped against time. In the mosaic, all such features may be shown as stacked or mapped against each other, since

they are different measurements in the time domain. As a result, such mosaic(s) may have the effect of creating interactions between any combination of the graphs in order to create more possibilities for the neural network to latch onto.

**[0043]** Mosaics may be generated with additional images including, but not limited to, a wavelet decomposition graph 550, a snippet of the waveform 560, a Voronoi diagram or other graphical representation of the reporting sensor array 570, a graph of recent nearby incidents 580, and/or certain features with respect to distance 590 where independent features may be mapped into the orthogonal color space of red, green and blue. The mosaic may also include one or more additional rows, such as a last row 595 including a visual representation of numeric metrics, the power spectrum, and a smaller version of the features with respect to distance 590B. Again, one objective is to get certain new features for the neural network to latch onto after the convolutions have blurred and distorted the image, so a repetition of cells is encouraged, as shown at the bottom of FIGs. 5A-5B. Figure 5B differs from Fig. 5A largely in having a *modified* graphical representation of the reporting sensor array, in field 599. Here, Fig. 5B illustrates a modified Voronoi diagram, at 599, which has been modified to remove features, such as graphical elements, of lesser, little or no importance to the classifier and/or classification, e.g., to yield improved processing. Here, for example, blue lines were removed from the subject Voronoi diagram to allow the classifier to avoid processing such feature(s), which improved performance. Of course, removal and/or modifications, like this, can be made to other graphical elements within all areas of the mosaics for elements where similar benefits may be achieved by such adjustments.

**[0044]** FIG. 6 is an illustration of another exemplary mosaic suitable for use in classification of acoustic events, consistent with certain aspects related to the innovations herein. The exemplary mosaic shown in FIG. 6 includes several features shown in the previous mosaic, i.e., FIG. 5. Referring to FIG. 6, a mosaic may include, at 610, a 2.0 second waveform signal frequency domain filter at a larger size than 510. Further, the mosaic may also include, at 620, a Voronoi diagram, such as one with a slightly different set of colors. Lastly, at 630, the mosaic may include the values of some pulse features in distance versus normalized value space. The representation shown in FIG. 6 is another example of a mosaic

that could be used, though it may be less effective than other implementations, for certain use cases, in terms of accuracy.

**[0045]** With regard to exemplary processing involving such a mosaic, from start to finish and referring to FIG. 1, sensor audio is used to determine a location in 110 and 120, then passed through a simple classifier at 130. Next, audio is downloaded from the sensors at 140, along with other non-acoustic data 150 to construct the mosaic, at 160. At this point, the mosaic image would be rendered like the examples shown in FIGs. 4-6, with audio from the first and second nearest audio waveform images, and then the images are passed to the Ensemble of Convolutional Neural Networks, at 170, to create an inference. Creation of such inferences occurs through many convolutions, which blurs many aspects of the image together. This blurring includes mixing colors together with 3x3 and 1x1 convolutions as well as areas of the image with 7x7 convolutions. As a result of such blurring and mixing, overlap and other similarities between differing sets of acoustic data can be advantageously analyzed including via the use of existing functionality already in place for image comparison. By leveraging the blurring of convolutional neural networks as well as a robust experience of the acoustic field, the combinations of this data in the form of a mosaic allow for analyzing data quickly and efficiently in real time. Not only do such combinations of data allow for automatic analysis of incoming data, they offer a unique perspective which is not able to be duplicated by any presently existing technology or means. Such unique perspective is achieved as a function of manipulating the unique combinations of the dataset, the convolutions used, and the composition order of the mosaic, any of which may be changed to create more advantageous, optimized and/or improved outputs or results, and which may also be selected or adjusted based on the data being analyzed and/or comparison(s) pursued.

**[0046]** FIG. 7 is an illustration of an exemplary system adapted to perform various features and/or functionality of the disclosed technology, consistent with certain aspects related to the innovations herein. As shown in FIG. 7, an exemplary system may include components associated with an event, incident or application 710, a processing system or platform 730, neural network(s) 740B that are external to the system or platform 730, other processing components 750 that are also external or remote from the system or platform 730, as well as one or more networks 720 through which any or all of such

computing components may be connected. With regard to an event, incident or application of interest, at 710, various sensors 712, other data sources 714, and/or processing components 716 may be utilized to obtain the information to be processing via the disclosed technology. As shown in FIG. 7, for example, a plurality of sensors 712A, 712B, 712C, which may be of different types, may be utilized to obtain different streams of raw data directly associated with the event, incident or application being analyzed. Further, various data sources other than the sensors 714 may be accessed or employed to yield additional information 715 regarding the event, incident or application of interest. Here, e.g. with regard to acoustic data processing implementations herein, this additional information 715 may include non-acoustic data such as a computed distance from the acoustic source to the receiver, the history of recent acoustic events in the same area as the event just detected, local elevation information, granular weather information, or any other data or information that may impact processing (identification, analysis, etc.) of the subject event, incident or application via the image mosaics herein. Lastly, all of the data and information from the sensors 712 and the other data sources 714 may be handled, aggregated and/or processed by one or more processors 716, as part of assembling and preparing this information either for transmission to the network(s) 720, at 718, and/or for transmission directly to the system or platform 730, at 719.

**[0047]** The event, incident or application data from 710 may then be processed at the system or platform 730. As shown in FIG. 7, an exemplary system or platform 730 may include one or more computer processors 732, one or more databases of mosaics 734, one or more neural networks 740A, as well as data regarding a subject mosaic design 736 that is used in the neural network processing described elsewhere herein. Further, such system or platform 730 may perform processing in connection with one or more external neural networks 740B and/or one or more other processing components 750 that may perform some, all and/or different processing as that described above in connection with the image mosaic processing of the disclosed technology.

**[0048]** Machine learning systems tend to overestimate the accuracy of their output predictions. However, if a monotonic relationship exists between the output of a machine classifier and actual classifier accuracy—if higher classifier output values imply higher real-world accuracy—then a map can be constructed between the two. One process for

calibrating classifier confidence is as follows: start with a set of input data (as potential gunfire incident images) that have been tagged by class; evaluate each incident image using a machine classification system and record the output confidence value; break the incident images into sets by output confidence value, with approximately the same number of incidents in each set; compute an estimate of the actual classifier accuracy for each set by computing the number of incidents correctly classified divided by the total number of incidents in the set; construct a map of the output of machine classifier confidence to classifier accuracy using the results for each set. This method provides a realistic output confidence for a classifier that can be used as a measure for comparing different kinds of classifiers, or comparing classifiers trained using different sets of data.

**[0049]** Such technique(s) can be used to increase classification reliability and/or confidence in an acoustic gunshot detection system. In a distributed acoustic sensor network that makes use of multilateration, sufficient sensor density is deployed such that a single acoustic event can be detected on two or more acoustic receivers. Each acoustic receiver will detect a different acoustic signal, a result of the asymmetry of the acoustic energy emitted by the firearm and the varying nature of the propagation path between the source and each individual receiver. In urban environments it is common for the direct acoustic path between source and receiver to be obstructed by buildings, resulting in a distorted signal that is more characteristic of the acoustic environment than of the acoustic source. Such distorted signals are more difficult for a machine-learning system to classify with high confidence. These issues can be overcome by measuring acoustic data on a plurality of acoustic receivers, generating an image mosaic for each receiver, classifying each mosaic independently using an appropriate classifier, mapping each output to a common confidence measure, and selecting the single classifier with the greatest normalized confidence as the best overall classification result. Using single-receiver images also improves the reliability of the image generation process, since successful download of audio data from a single receiver from the plurality of acoustic receivers on which downloads are attempted is sufficient to construct a valid image for classification, as is shown by the following calculation: if the probability of audio an download failure is 1% and independent of other factors and downloads are attempted from the two nearest sensors, then the probability of a successful download from both sensors is 98.01% while the probability of a successful download from either sensor is

99.99%. Creating independent image mosaics from multiple sensors results in higher classifier accuracy when multiple images are generated and fewer cases where no acoustic data is available for classification.

**[0050]** An alternative implementation for combining data from multiple receivers is to build a mosaic image containing acoustic data from all acoustic sensors using the methods described in the disclosed technology. This potentially increases classifier accuracy by providing more input data to a machine learning system. The downside of using acoustic data from multiple receivers in a single image is that it requires a considerably larger amount of training data (up to the square of the number of images required with a single acoustic receiver per image) to achieve the same level of training, and increases the likelihood that one or more of the acoustic waveforms necessary for image mosaic construction will not be available.

**[0051]** As disclosed herein, implementations and features of the present inventions may be implemented through computer-hardware, software and/or firmware. For example, the systems and methods disclosed herein may be embodied in various forms including, for example, one or more data processors, such as computer(s), server(s) and the like, and may also include or access at least one database, digital electronic circuitry, firmware, software, or in combinations of them. Further, while some of the disclosed implementations describe specific (e.g., hardware, etc.) components, systems and methods consistent with the innovations herein may be implemented with any combination of hardware, software and/or firmware. Moreover, the above-noted features and other aspects and principles of the innovations herein may be implemented in various environments. Such environments and related applications may be specially constructed for performing the various processes and operations according to the inventions or they may include a general-purpose computer or computing platform selectively activated or reconfigured by code to provide the necessary functionality. The processes disclosed herein are not inherently related to any particular computer, network, architecture, environment, or other apparatus, and may be implemented by a suitable combination of hardware, software, and/or firmware. For example, various general-purpose machines may be used with programs written in accordance with teachings of the inventions, or it may be more convenient to construct a specialized apparatus or system to perform the required methods and techniques.

**[0052]** In the present description, the terms component, module, device, etc. may refer to any type of logical or functional device, process or blocks that may be implemented in a variety of ways. For example, the functions of various blocks can be combined with one another and/or distributed into any other number of modules. Each module can be implemented as a software program stored on a tangible memory (e.g., random access memory, read only memory, CD-ROM memory, hard disk drive) within or associated with the computing elements, sensors, receivers, etc. disclosed above, e.g., to be read by a processing unit to implement the functions of the innovations herein. Also, the modules can be implemented as hardware logic circuitry implementing the functions encompassed by the innovations herein. Finally, the modules can be implemented using special purpose instructions (SIMD instructions), field programmable logic arrays or any mix thereof which provides the desired level performance and cost.

**[0053]** Aspects of the systems and methods described herein may be implemented as functionality programmed into any of a variety of circuitry, including programmable logic devices (PLDs), such as field programmable gate arrays (FPGAs), programmable array logic (PAL) devices, electrically programmable logic and memory devices and standard cell-based devices, as well as application specific integrated circuits. Some other possibilities for implementing aspects include: memory devices, microcontrollers with memory (such as EEPROM), embedded microprocessors, firmware, software, etc. Furthermore, aspects may be embodied in microprocessors having software-based circuit emulation, discrete logic (sequential and combinatorial), custom devices, fuzzy logic, neural networks, other AI (Artificial Intelligence) or machine learning systems, quantum devices, and hybrids of any of the above device types.

**[0054]** It should also be noted that various logic and/or features disclosed herein may be enabled using any number of combinations of hardware, firmware, and/or as data and/or instructions embodied in various machine-readable or computer-readable media, in terms of their behavioral, register transfer, logic component, and/or other characteristics. Computer-readable media in which such formatted data and/or instructions may be embodied include, but are not limited to, non-volatile storage media in tangible various forms (e.g., optical, magnetic or semiconductor storage media), though do not encompass transitory media.

**[0055]** Unless the context clearly requires otherwise, throughout the description, the words "comprise," "comprising," and the like are to be construed in an inclusive sense as opposed to an exclusive or exhaustive sense; that is to say, in a sense of "including, but not limited to." Words using the singular or plural number also include the plural or singular number respectively. Additionally, the words "herein," "hereunder," "above," "below," and words of similar import refer to this application as a whole and not to any particular portions of this application. When the word "or" is used in reference to a list of two or more items, that word covers all of the following interpretations of the word: any of the items in the list, all of the items in the list and any combination of the items in the list.

**[0056]** Other implementations of the inventions will be apparent to those skilled in the art from consideration of the specification and practice of the innovations disclosed herein. It is intended that the specification and examples be considered as exemplary only, with a true scope and spirit of the inventions being indicated by the present disclosure and claims and various associated principles of related patent doctrine.

## Claims:

1. A method for constructing image mosaics suitable for machine classification of an acoustic event, the image mosaics comprised of acoustic data, the method comprising:
  - generating at least one first graphical representation of one or more acoustic signals associated with the acoustic data;
  - generating at least one second graphical representation of non-acoustic signals associated with the non-acoustic data;
  - assembling one or more elements of the first graphical representation and the second graphical representation to generate a single image mosaic; and
  - utilizing the single image mosaic to train a machine learning engine for the purpose of classifying acoustic events.
2. The method of claim 1 or the invention of any claim herein, wherein the image mosaics further comprise non-acoustic data.
3. The method of claim 1 or the invention of any claim herein, wherein the machine learning engine comprises a convolutional neural network that uses the image mosaics to provide improved matching of audio impulses with potential gunshot events by constructing the image mosaics with image subcomponents that are arranged in a format wherein the image subcomponents are positioned in fixed/spaced relation to each other and/or have their domains aligned, such that the format is utilized by the convolutional layers to maximize recognition of common features achieved via convolutional processing.
4. The method of claim 3 or the invention of any claim herein, wherein the format includes aligning the domains of similar and different measurements together.
5. The method of claims 1 or 3 or the invention of any claim herein, wherein two or more graphical representations (subcomponents) are combined into a same image mosaic, wherein the subcomponents are positioned in spaced relation (e.g., adjacent etc) to each other, arranged, aligned, and/or repeated such that comparison or recognition of the image mosaic

against a master mosaic image is improved by comparison of adjusted (e.g., blurred, distorted, etc.) images after convolutional operations performed in the convolutional layers of the neural network.

6. The method of claim 3 or 5 or the invention of any claim herein, wherein the subcomponents of the image are placed in predefined compositional order (location/position/arrangement within the image) that is consistent with the location/position/arrangement of like subcomponents in a master mosaic image to facilitate image machine learning and especially convolutions.

7. The method of claim 1 or the invention of any claim herein, wherein image-based machine learning classifiers are combined with non-acoustic data, such as computed distance from the acoustic source to the receiver, history of recent acoustic events in the same area as the event just detected, and/or other non-acoustic data associated with the acoustic source, event, etc.,

8. The method of claim 1 or the invention of any claim herein, wherein the image mosaics comprise two or more image subcomponents positioned in fixed spatial relationship to each other and two or more graphical subcomponents aligned with respect to each other along an axis characterized by a common domain shared by the graphical subcomponents.

9. The method of claim 1 or the invention of any claim herein, wherein the image mosaics comprises:

a first field or row comprised of two or more graphical subcomponents aligned with respect to each other along an axis characterized by a common domain shared by the graphical subcomponents;

a second field or row comprised of two or more image subcomponents positioned in fixed spatial relationship to each other; and

a third field or row comprised of two or more image subcomponents positioned in fixed spatial relationship to each other, wherein the image components of the third field or

row are arranged and/or repeated to emphasize recognition via convolutional neural network processing.

10. The method of claim 9 or the invention of any claim herein, wherein the common domain comprises at least one of timeframe, frequency, area, distance-frequency, distance-timeframe, volume, power, or energy.

11. The method of claim 1 or the invention of any claim herein, wherein the image subcomponents are arranged in combinations that more-readily identify acoustic characteristics common between a new mosaic (an acoustic event to be processed) and known mosaics (known acoustic events, such as known gunshot impulses) in that the combinations yield adjusted/recognizable (e.g., blurred, distorted, etc.) graphical results that depict matching event characteristics via generation of the graphical results using convolutional neural network processing.

12. The method of claim 1 or claim 2 or the invention of any claim herein, wherein the single image mosaic further comprises a combination of waveform data from multiple transducers, such as microphones, accelerometers, infrared detectors or signals, optical images, or chemical sensors, into a single image.

13. A method for combining waveform data regarding an event received from multiple transducers into a single image mosaic, the multiple transducers including at least one of a microphone, an accelerometer, an infrared detector, an optical sensor, and/or a chemical sensor, etc., the method comprising:

generating at least one first graphical representation of one or more acoustic signals associated with the acoustic data;

generating at least one second graphical representation of non-acoustic signals associated with the non-acoustic data;

assembling one or more elements of the first graphical representation and the second graphical representation to generate the single image mosaic having a specified format;

wherein the single image mosaic is constructed and configured for processing by a machine learning image classifier to classify the acoustic event against mosaic images having the specified format.

14. The method of claim 13 or the invention of any claim herein, wherein the single image mosaic is comprised of subcomponents associated with an acoustic event, wherein the subcomponents are positioned in spaced relation to each other, arranged, aligned, and/or repeated such that comparison or recognition of the image mosaic against a master mosaic image is improved by comparison of blurred and distorted images after convolutional operations performed in the convolutional layers of the neural network.

15. A method for constructing image mosaics comprised of a plurality of subcomponents that are graphical indicia associated with an acoustic source/event, wherein the subcomponents are positioned in spaced relation to each other, arranged, aligned, and/or repeated such that any data associated with the acoustic source/event is suitable for processing by a machine learning system that performs image processing, such as performing image recognition of the image mosaics, via a convolution neural network, to provide data processing of the acoustic source/event without processing of underlying acoustic data.

16. A method for increasing the reliability and accuracy of machine classification of acoustic point source incidents, the method comprising:

generating image mosaics for use in the machine classification, the image mosaics comprised of a plurality of image subcomponents that are created based on information associated with an acoustic point source incident; and

processing acoustic information by:

performing processing (e.g., generating requests, etc.) associated with a plurality of recording devices to obtain signals derived from the recording devices;

generating the image subcomponents based on the signals derived from the recording devices; and

aligning, within each of the mosaic images, the image subcomponents (based on the signals derived from the recording devices) in the time domain such that the acoustic incident appears in a same part/spatial position of each image mosaic.

17. A method for improving reliability and/or quality of a classifier system, the method comprising:

training multiple models using one or more machine learning systems by generating image mosaics comprised of image subcomponents that are created based on information associated with acoustic events/sources;

applying the image mosaic to all trained systems; and

analyzing output result from all responding machine learning systems, so as to adjust to an actionable probability from the skewed output, and then combine the responses for an increased quality in classification or an increase in reliability with two (or two or more, or one or more) possible images returning.

18. A method for processing image mosaics suitable for machine classification of non-image data, the image mosaics comprised of image data and other graphical indicia derived based on the non-image data, the method comprising:

generating the image mosaics for use in the machine classification, the image mosaics comprised of a plurality of image subcomponents that are created based on information associated with an incident or event characterized by the non-image data;

assembling and/or arranging the data in a spatial orientation within a single image mosaic as a function of correlating factors associated with the non-image data underlying the image mosaics, the correlating factors including data that enables a test incident or event to be correlated with (matched against) known incidents or events; and

analyzing an image mosaic of the test incident or event against the image mosaics of the known incidents or events, wherein recognition of the test incident or event is improved by comparison of processed (e.g., blurred, distorted, etc.) images derived via neural network, CNN (convolutional neural network) and/or other such machine classification operations.

19. A method for combining waveform data regarding an acoustic event received from multiple transducers into a single image mosaic, the method comprising:

generating at least one first graphical representation of one or more acoustic signals associated with the acoustic event;

generating at least one second graphical representation of non-acoustic data associated with the acoustic event;

assembling elements of the first graphical representation and the second graphical representation to generate the single image mosaic having a specified format;

wherein the single image mosaic is constructed and configured for processing by a machine learning image classifier to classify the acoustic event via processing of the single mosaic image against mosaic images having the specified format.

20. The method of claim 19 or the invention of any claim herein wherein the mosaic images against which the single mosaic image is processed comprise known acoustic events.

21. The method of claim 1 or the invention of any claim herein, wherein a plurality of image mosaics is constructed using acoustic data captured at different geographic locations and classified either together or independently, in order to increase the confidence in the output of the machine learning system.

22. The method of claim 1 or the invention of any claim herein, wherein a plurality of image mosaics is constructed using acoustic data captured at different geographic locations and classified independently, which may, in some implementations, increase the confidence and/or reliability in the output of the machine learning system.

23. The method of claim 1 or the invention of any claim herein, wherein a plurality of image mosaics is constructed using acoustic data captured at different geographic locations and classified either together or independently, which may, in some implementations, increase classifier accuracy by providing more input data to a machine learning system.

24. A system for processing impulsive acoustic events (e.g., locating and/or classifying gunshots, etc.), the system comprising:

one or more servers, computer processors, and/or computer readable media configured to:

receive acoustic and/or other data; and

perform one or more portions or steps of any of claims 1-23 and/or other features or functionality set forth elsewhere in the present disclosure.

25. One or more computer-readable media containing and/or configured to execute computer-readable instructions, the computer-readable instructions comprising instructions that, when executed by one or more processors, cause the one or more processors to:

perform one or more portions, aspects and/or steps of any of claims 1-24 and/or other features or functionality set forth elsewhere in this disclosure.

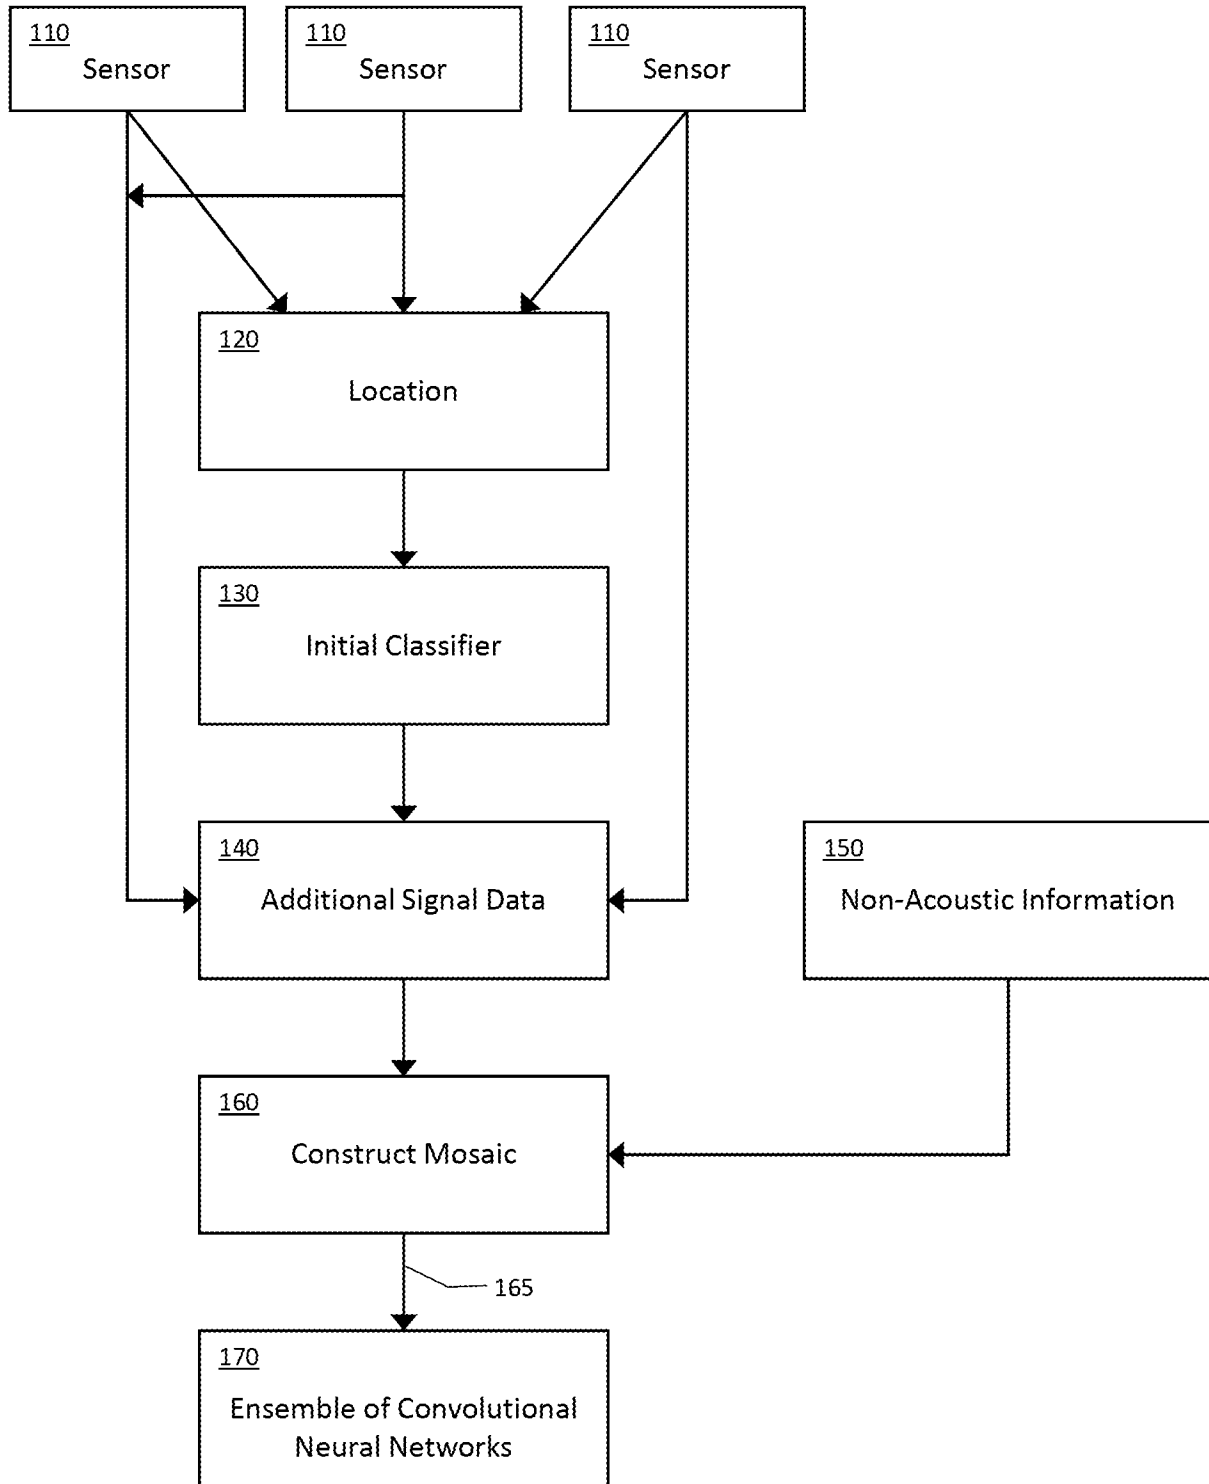


FIG. 1


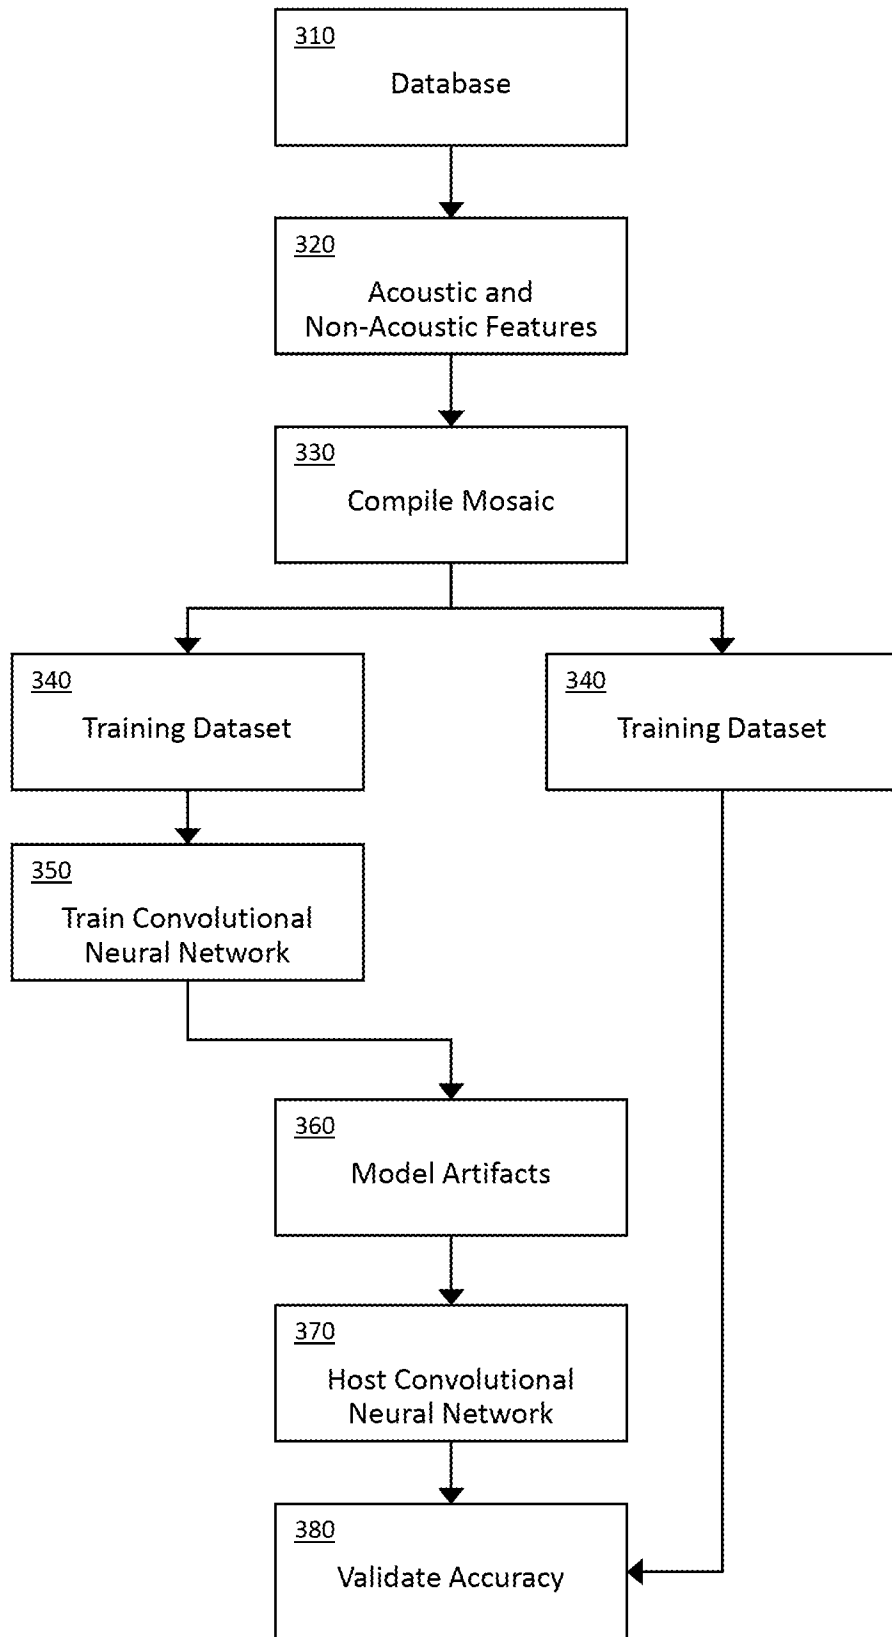


FIG. 2



**FIG. 3**




FIG. 4

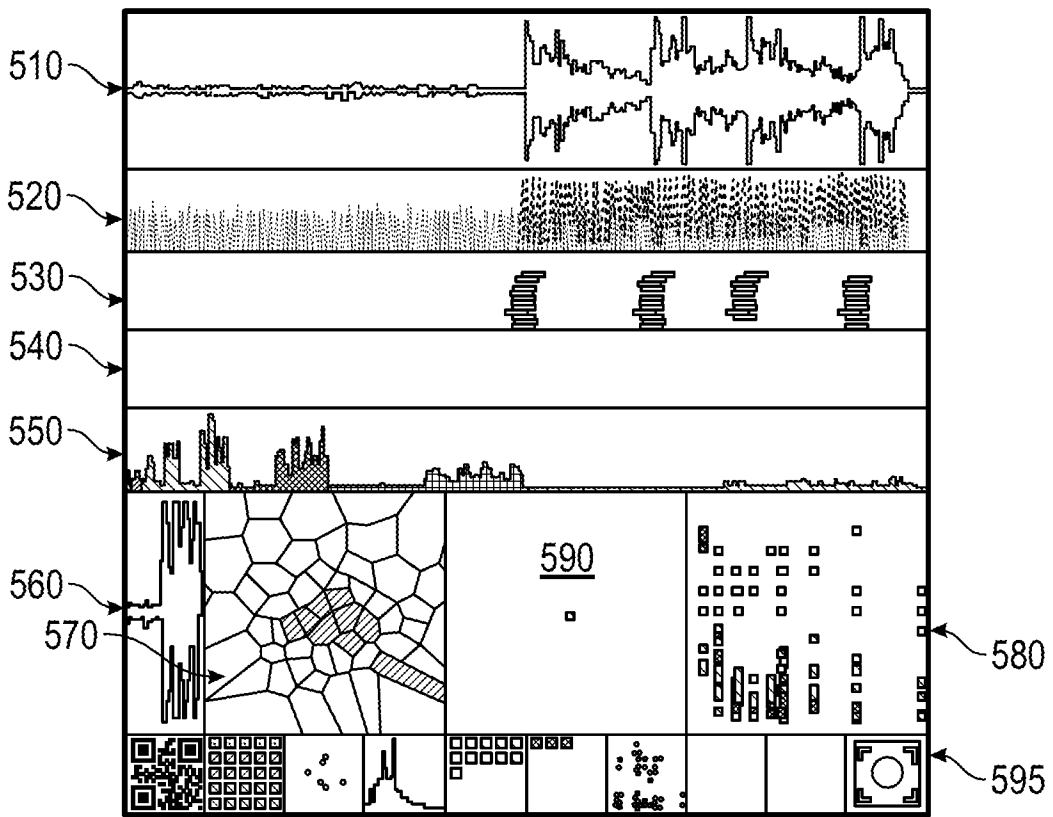


FIG. 5A

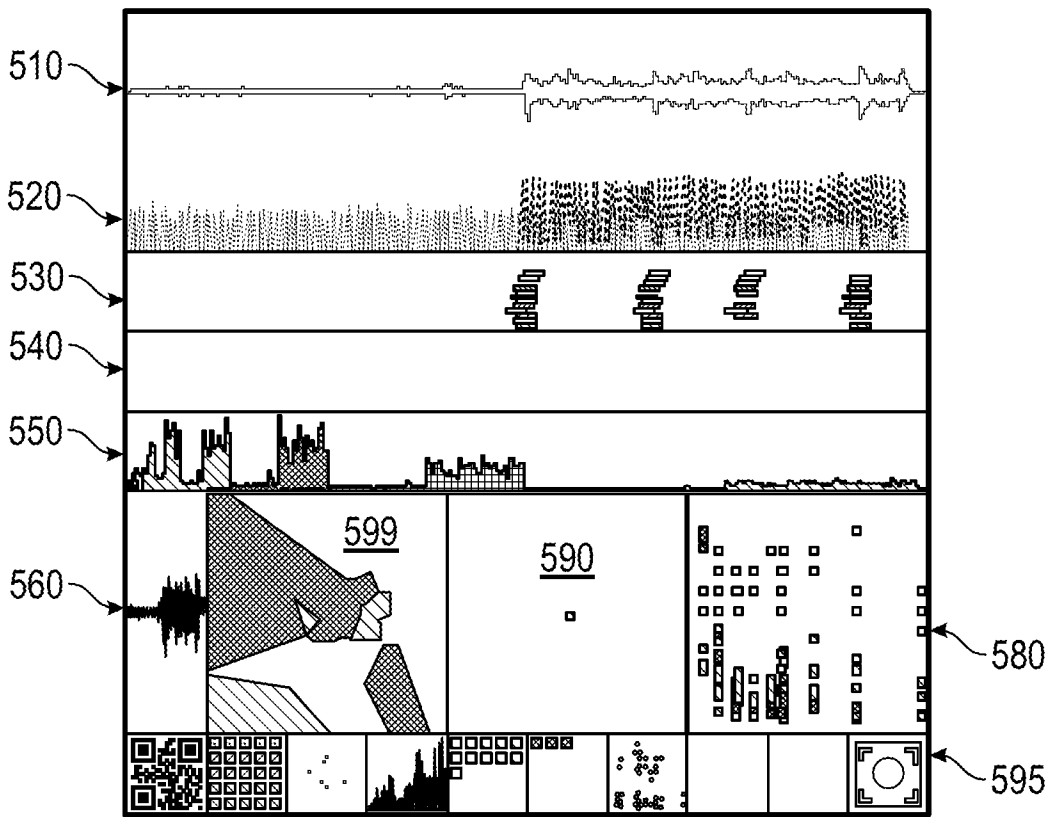


FIG. 5B




FIG. 6




FIG. 7

## INTERNATIONAL SEARCH REPORT

International application No.

PCT/US 2020/018697

A. CLASSIFICATION OF SUBJECT MATTER		
<i>G06K 9/62 (2006.01)</i> <i>G06N 20/00 (2006.01)</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)		
G06F 17/15, G06K 9/62, G06N 20/00, 3/04, G06T 3/40, 5/00		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)		
PatSearch (RUPTO internal), USPTO, PAJ, Esp@cenet, DWPI, EAPATIS, PATENTSCOPE		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2016277863 A1 (INTEL CORP) 22.09.2016	1-4, 7-11, 13-25
A	CN 102282571 A (NOKIA CORP) 14.12.2011	1-4, 7-11, 13-25
A	US 2018260626 A1 (ACCENTURE GLOBAL SERVICES LTD) 13.09.2018	1-4, 7-11, 13-25
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> 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 document but published on or after the international filing date  "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)  "O" document referring to an oral disclosure, use, exhibition or other means  "P" document published prior to the international filing date but later than the priority date claimed	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention  "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone  "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art  "&" document member of the same patent family	
Date of the actual completion of the international search		Date of mailing of the international search report
19 May 2020 (19.05.2020)		28 May 2020 (28.05.2020)
Name and mailing address of the ISA/RU: Federal Institute of Industrial Property, Berezhkovskaya nab., 30-1, Moscow, G-59, GSP-3, Russia, 125993 Facsimile No: (8-495) 531-63-18, (8-499) 243-33-37		Authorized officer  M. Markov  Telephone No. 8 499 240 25 91

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

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

1.  Claims Nos.:  
because they relate to subject matter not required to be searched by this Authority, namely:
  
2.  Claims Nos.:  
because they relate to parts of the international application that do not comply with the prescribed requirements to such an extent that no meaningful international search can be carried out, specifically:
  
3.  Claims Nos.: 5-6, 12  
because they are dependent claims and are not drafted in accordance with the second and third sentences of Rule 6.4(a).

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

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

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

**Remark on Protest**

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