

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

(19) World Intellectual Property Organization
International Bureau



(10) International Publication Number

WO 2015/123460 A1

(43) International Publication Date
20 August 2015 (20.08.2015)

(51) International Patent Classification:
G06N 3/04 (2006.01) *G10L 25/30* (2013.01)

(74) Agents: LENKIN, Alan M. et al.; Seyfarth Shaw LLP, Suite 3500, 2029 Century Park East, Los Angeles, California 90067-3021 (US).

(21) International Application Number:

PCT/US2015/015697

(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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(22) International Filing Date:
12 February 2015 (12.02.2015)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
61/940,281 14 February 2014 (14.02.2014) US
14/286,556 23 May 2014 (23.05.2014) US

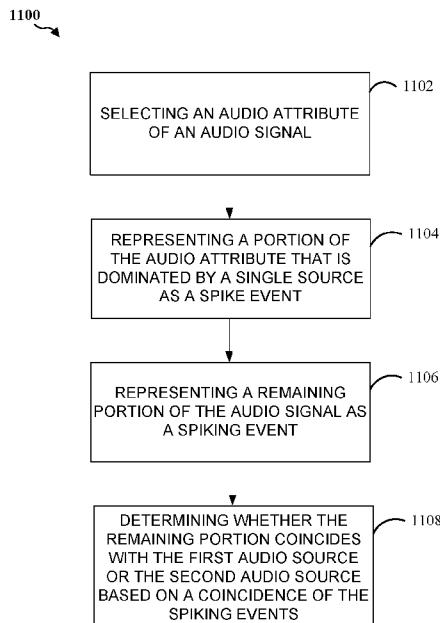
(71) Applicant: QUALCOMM INCORPORATED [US/US];
ATTN: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

(72) Inventors: KRISHNAN, Lakshmi; 3407 Tulane Drive, Apt. 33, Hyattsville, Maryland 20783 (US). CAMPOS, Michael; 5775 Morehouse Drive, San Diego, California 92121-1714 (US).

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

[Continued on next page]

(54) Title: AUDITORY SOURCE SEPARATION IN A SPIKING NEURAL NETWORK



(57) Abstract: A method of audio source segregation includes selecting an audio attribute of an audio signal. The method also includes representing a portion of the audio attribute that is dominated by a single source as a source spiking event. In addition, the method includes representing a remaining portion of the audio signal as an audio signal spiking event. The method further includes determining whether the remaining portion coincides with the single source based on coincidence of the source spiking event and audio signal spiking event.

FIG. 11



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

as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

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

AUDITORY SOURCE SEPARATION IN A SPIKING NEURAL NETWORK

CROSS-REFERENCE TO RELATED APPLICATION

[0001] The present application claims the benefit of U.S. Provisional Patent Application No. 61/940,281, filed on February 14, 2014, and titled “AUDITORY SOURCE SEPARATION IN A SPIKING NEURAL NETWORK,” the disclosure of which is expressly incorporated by reference herein in its entirety.

BACKGROUND

Field

[0002] Certain aspects of the present disclosure generally relate to neural system engineering and, more particularly, to systems and methods for auditory source separation in a spiking neural network.

Background

[0003] An artificial neural network, which may comprise an interconnected group of artificial neurons (i.e., neuron models), is a computational device or represents a method to be performed by a computational device. Artificial neural networks may have corresponding structure and/or function in biological neural networks. However, artificial neural networks may provide innovative and useful computational techniques for certain applications in which traditional computational techniques are cumbersome, impractical, or inadequate. Because artificial neural networks can infer a function from observations, such networks are particularly useful in applications where the complexity of the task or data makes the design of the function by conventional techniques burdensome.

SUMMARY

[0004] In an aspect of the present disclosure, a method of audio source segregation is disclosed. The method includes selecting an audio attribute of an audio signal. The method also includes representing a portion of the audio attribute that is dominated by a single source as a source spiking event. In addition, the method includes representing a remaining portion of the audio signal as an audio signal spiking event. The method

further includes determining whether the remaining portion coincides with the single source based on coincidence of the source spiking event and audio signal spiking event.

[0005] In another aspect of the present disclosure, an apparatus for audio source segregation is disclosed. The apparatus has a memory and at least one processor. The processor(s) is coupled to the memory and is configured to select an audio attribute of an audio signal. The processor(s) is also configured to represent a portion of the audio attribute that is dominated by a single source as a source spiking event. In addition, the processor(s) is configured to represent a remaining portion of the audio signal as an audio signal spiking event. The processor(s) is further configured to determine whether the remaining portion coincides with the single source based on coincidence of the source spiking event and audio signal spiking event.

[0006] In yet another aspect of the present disclosure, an apparatus for audio source segregation is disclosed. The apparatus includes means for selecting an audio attribute of an audio signal. The apparatus also includes means for representing a portion of the audio attribute that is dominated by a single source as a source spiking event. The apparatus additionally includes means for representing a remaining portion of the audio signal as an audio signal spiking event. The apparatus further includes means for determining whether the remaining portion coincides with the single source based on coincidence of the source spiking event and audio signal spiking event.

[0007] In still yet another aspect of the present disclosure, a computer program product for audio source segregation is disclosed. The computer program product has a non-transitory computer readable medium having encoded thereon program code. The program code includes program code to select an audio attribute of an audio signal. The program code also includes program code to represent a portion of the audio attribute that is dominated by a single source as a source spiking event. In addition, the program code includes program code to represent a remaining portion of the audio signal as an audio signal spiking event. The program code further includes program code to determine whether the remaining portion coincides with the single source based on coincidence of the source spiking event and audio signal spiking event.

[0008] This has outlined, rather broadly, the features and technical advantages of the present disclosure in order that the detailed description that follows may be better

understood. Additional features and advantages of the disclosure will be described below. It should be appreciated by those skilled in the art that this disclosure may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present disclosure. It should also be realized by those skilled in the art that such equivalent constructions do not depart from the teachings of the disclosure as set forth in the appended claims. The novel features, which are believed to be characteristic of the disclosure, both as to its organization and method of operation, together with further objects and advantages, will be better understood from the following description when considered in connection with the accompanying figures. It is to be expressly understood, however, that each of the figures is provided for the purpose of illustration and description only and is not intended as a definition of the limits of the present disclosure.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The features, nature, and advantages of the present disclosure will become more apparent from the detailed description set forth below when taken in conjunction with the drawings in which like reference characters identify correspondingly throughout.

[0010] FIGURE 1 illustrates an example network of neurons in accordance with certain aspects of the present disclosure.

[0011] FIGURE 2 illustrates an example of a processing unit (neuron) of a computational network (neural system or neural network) in accordance with certain aspects of the present disclosure.

[0012] FIGURE 3 illustrates an example of spike-timing dependent plasticity (STDP) curve in accordance with certain aspects of the present disclosure.

[0013] FIGURE 4 illustrates an example of a positive regime and a negative regime for defining behavior of a neuron model in accordance with certain aspects of the present disclosure.

[0014] FIGURE 5 illustrates an example implementation of designing a neural network using a general-purpose processor in accordance with certain aspects of the present disclosure.

[0015] FIGURE 6 illustrates an example implementation of designing a neural network where a memory may be interfaced with individual distributed processing units in accordance with certain aspects of the present disclosure.

[0016] FIGURE 7 illustrates an example implementation of designing a neural network based on distributed memories and distributed processing units in accordance with certain aspects of the present disclosure.

[0017] FIGURE 8 illustrates an example implementation of a neural network in accordance with certain aspects of the present disclosure.

[0018] FIGURES 9 and 10 are block diagrams illustrating exemplary architectures for auditory source segregation in accordance with aspects of the present disclosure.

[0019] FIGURE 11 is a block diagram illustrating a method for auditory source segregation in accordance with an aspect of the present disclosure.

DETAILED DESCRIPTION

[0020] The detailed description set forth below, in connection with the appended drawings, is intended as a description of various configurations and is not intended to represent the only configurations in which the concepts described herein may be practiced. The detailed description includes specific details for the purpose of providing a thorough understanding of the various concepts. However, it will be apparent to those skilled in the art that these concepts may be practiced without these specific details. In some instances, well-known structures and components are shown in block diagram form in order to avoid obscuring such concepts.

[0021] Based on the teachings, one skilled in the art should appreciate that the scope of the disclosure is intended to cover any aspect of the disclosure, whether implemented independently of or combined with any other aspect of the disclosure. For example, an apparatus may be implemented or a method may be practiced using any

number of the aspects set forth. In addition, the scope of the disclosure is intended to cover such an apparatus or method practiced using other structure, functionality, or structure and functionality in addition to or other than the various aspects of the disclosure set forth. It should be understood that any aspect of the disclosure disclosed may be embodied by one or more elements of a claim.

[0022] The word “exemplary” is used herein to mean “serving as an example, instance, or illustration.” Any aspect described herein as “exemplary” is not necessarily to be construed as preferred or advantageous over other aspects.

[0023] Although particular aspects are described herein, many variations and permutations of these aspects fall within the scope of the disclosure. Although some benefits and advantages of the preferred aspects are mentioned, the scope of the disclosure is not intended to be limited to particular benefits, uses or objectives. Rather, aspects of the disclosure are intended to be broadly applicable to different technologies, system configurations, networks and protocols, some of which are illustrated by way of example in the figures and in the following description of the preferred aspects. The detailed description and drawings are merely illustrative of the disclosure rather than limiting, the scope of the disclosure being defined by the appended claims and equivalents thereof.

AN EXAMPLE NEURAL SYSTEM, TRAINING AND OPERATION

[0024] FIGURE 1 illustrates an example artificial neural system 100 with multiple levels of neurons in accordance with certain aspects of the present disclosure. The neural system 100 may have a level of neurons 102 connected to another level of neurons 106 through a network of synaptic connections 104 (i.e., feed-forward connections). For simplicity, only two levels of neurons are illustrated in FIGURE 1, although fewer or more levels of neurons may exist in a neural system. It should be noted that some of the neurons may connect to other neurons of the same layer through lateral connections. Furthermore, some of the neurons may connect back to a neuron of a previous layer through feedback connections.

[0025] As illustrated in FIGURE 1, each neuron in the level 102 may receive an input signal 108 that may be generated by neurons of a previous level (not shown in FIGURE 1). The signal 108 may represent an input current of the level 102 neuron.

This current may be accumulated on the neuron membrane to charge a membrane potential. When the membrane potential reaches its threshold value, the neuron may fire and generate an output spike to be transferred to the next level of neurons (e.g., the level 106). In some modeling approaches, the neuron may continuously transfer a signal to the next level of neurons. This signal is typically a function of the membrane potential. Such behavior can be emulated or simulated in hardware and/or software, including analog and digital implementations such as those described below.

[0026] In biological neurons, the output spike generated when a neuron fires is referred to as an action potential. This electrical signal is a relatively rapid, transient, nerve impulse, having an amplitude of roughly 100 mV and a duration of about 1 ms. In a particular embodiment of a neural system having a series of connected neurons (e.g., the transfer of spikes from one level of neurons to another in FIGURE 1), every action potential has basically the same amplitude and duration, and thus, the information in the signal may be represented only by the frequency and number of spikes, or the time of spikes, rather than by the amplitude. The information carried by an action potential may be determined by the spike, the neuron that spiked, and the time of the spike relative to other spike or spikes. The importance of the spike may be determined by a weight applied to a connection between neurons, as explained below.

[0027] The transfer of spikes from one level of neurons to another may be achieved through the network of synaptic connections (or simply “synapses”) 104, as illustrated in FIGURE 1. Relative to the synapses 104, neurons of level 102 may be considered presynaptic neurons and neurons of level 106 may be considered postsynaptic neurons. The synapses 104 may receive output signals (i.e., spikes) from the level 102 neurons and scale those signals according to adjustable synaptic weights $w_1^{(i,i+1)}, \dots, w_P^{(i,i+1)}$ where P is a total number of synaptic connections between the neurons of levels 102 and 106 and i is an indicator of the neuron level. In the example of FIGURE 1, i represents neuron level 102 and $i+1$ represents neuron level 106. Further, the scaled signals may be combined as an input signal of each neuron in the level 106. Every neuron in the level 106 may generate output spikes 110 based on the corresponding combined input signal. The output spikes 110 may be transferred to another level of neurons using another network of synaptic connections (not shown in FIGURE 1).

[0028] Biological synapses can mediate either excitatory or inhibitory (hyperpolarizing) actions in postsynaptic neurons and can also serve to amplify neuronal signals. Excitatory signals depolarize the membrane potential (i.e., increase the membrane potential with respect to the resting potential). If enough excitatory signals are received within a certain time period to depolarize the membrane potential above a threshold, an action potential occurs in the postsynaptic neuron. In contrast, inhibitory signals generally hyperpolarize (i.e., lower) the membrane potential. Inhibitory signals, if strong enough, can counteract the sum of excitatory signals and prevent the membrane potential from reaching a threshold. In addition to counteracting synaptic excitation, synaptic inhibition can exert powerful control over spontaneously active neurons. A spontaneously active neuron refers to a neuron that spikes without further input, for example due to its dynamics or a feedback. By suppressing the spontaneous generation of action potentials in these neurons, synaptic inhibition can shape the pattern of firing in a neuron, which is generally referred to as sculpturing. The various synapses 104 may act as any combination of excitatory or inhibitory synapses, depending on the behavior desired.

[0029] The neural system 100 may be emulated by a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device (PLD), discrete gate or transistor logic, discrete hardware components, a software module executed by a processor, or any combination thereof. The neural system 100 may be utilized in a large range of applications, such as image and pattern recognition, machine learning, motor control, and alike. Each neuron in the neural system 100 may be implemented as a neuron circuit. The neuron membrane charged to the threshold value initiating the output spike may be implemented, for example, as a capacitor that integrates an electrical current flowing through it.

[0030] In an aspect, the capacitor may be eliminated as the electrical current integrating device of the neuron circuit, and a smaller memristor element may be used in its place. This approach may be applied in neuron circuits, as well as in various other applications where bulky capacitors are utilized as electrical current integrators. In addition, each of the synapses 104 may be implemented based on a memristor element, where synaptic weight changes may relate to changes of the memristor resistance. With nanometer feature-sized memristors, the area of a neuron circuit and synapses may be

substantially reduced, which may make implementation of a large-scale neural system hardware implementation more practical.

[0031] Functionality of a neural processor that emulates the neural system 100 may depend on weights of synaptic connections, which may control strengths of connections between neurons. The synaptic weights may be stored in a non-volatile memory in order to preserve functionality of the processor after being powered down. In an aspect, the synaptic weight memory may be implemented on a separate external chip from the main neural processor chip. The synaptic weight memory may be packaged separately from the neural processor chip as a replaceable memory card. This may provide diverse functionalities to the neural processor, where a particular functionality may be based on synaptic weights stored in a memory card currently attached to the neural processor.

[0032] FIGURE 2 illustrates an exemplary diagram 200 of a processing unit (e.g., a neuron or neuron circuit) 202 of a computational network (e.g., a neural system or a neural network) in accordance with certain aspects of the present disclosure. For example, the neuron 202 may correspond to any of the neurons of levels 102 and 106 from FIGURE 1. The neuron 202 may receive multiple input signals 204₁-204_N, which may be signals external to the neural system, or signals generated by other neurons of the same neural system, or both. The input signal may be a current, a conductance, a voltage, a real-valued, and/or a complex-valued. The input signal may comprise a numerical value with a fixed-point or a floating-point representation. These input signals may be delivered to the neuron 202 through synaptic connections that scale the signals according to adjustable synaptic weights 206₁-206_N (W₁-W_N), where N may be a total number of input connections of the neuron 202.

[0033] The neuron 202 may combine the scaled input signals and use the combined scaled inputs to generate an output signal 208 (i.e., a signal Y). The output signal 208 may be a current, a conductance, a voltage, a real-valued and/or a complex-valued. The output signal may be a numerical value with a fixed-point or a floating-point representation. The output signal 208 may be then transferred as an input signal to other neurons of the same neural system, or as an input signal to the same neuron 202, or as an output of the neural system.

[0034] The processing unit (neuron) 202 may be emulated by an electrical circuit, and its input and output connections may be emulated by electrical connections with synaptic circuits. The processing unit 202 and its input and output connections may also be emulated by a software code. The processing unit 202 may also be emulated by an electric circuit, whereas its input and output connections may be emulated by a software code. In an aspect, the processing unit 202 in the computational network may be an analog electrical circuit. In another aspect, the processing unit 202 may be a digital electrical circuit. In yet another aspect, the processing unit 202 may be a mixed-signal electrical circuit with both analog and digital components. The computational network may include processing units in any of the aforementioned forms. The computational network (neural system or neural network) using such processing units may be utilized in a large range of applications, such as image and pattern recognition, machine learning, motor control, and the like.

[0035] During the course of training a neural network, synaptic weights (e.g., the weights $w_1^{(i,i+1)}, \dots, w_P^{(i,i+1)}$ from FIGURE 1 and/or the weights 206₁-206_N from FIGURE 2) may be initialized with random values and increased or decreased according to a learning rule. Those skilled in the art will appreciate that examples of the learning rule include, but are not limited to the spike-timing-dependent plasticity (STDP) learning rule, the Hebb rule, the Oja rule, the Bienenstock-Copper-Munro (BCM) rule, etc. In certain aspects, the weights may settle or converge to one of two values (i.e., a bimodal distribution of weights). This effect can be utilized to reduce the number of bits for each synaptic weight, increase the speed of reading and writing from/to a memory storing the synaptic weights, and to reduce power and/or processor consumption of the synaptic memory.

Synapse Type

[0036] In hardware and software models of neural networks, the processing of synapse related functions can be based on synaptic type. Synapse types may be non-plastic synapses (no changes of weight and delay), plastic synapses (weight may change), structural delay plastic synapses (weight and delay may change), fully plastic synapses (weight, delay and connectivity may change), and variations thereupon (e.g., delay may change, but no change in weight or connectivity). The advantage of multiple types is that processing can be subdivided. For example, non-plastic synapses may not

use plasticity functions to be executed (or waiting for such functions to complete). Similarly, delay and weight plasticity may be subdivided into operations that may operate together or separately, in sequence or in parallel. Different types of synapses may have different lookup tables or formulas and parameters for each of the different plasticity types that apply. Thus, the methods would access the relevant tables, formulas, or parameters for the synapse's type.

[0037] There are further implications of the fact that spike-timing dependent structural plasticity may be executed independently of synaptic plasticity. Structural plasticity may be executed even if there is no change to weight magnitude (e.g., if the weight has reached a minimum or maximum value, or it is not changed due to some other reason) s structural plasticity (i.e., an amount of delay change) may be a direct function of pre-post spike time difference. Alternatively, structural plasticity may be set as a function of the weight change amount or based on conditions relating to bounds of the weights or weight changes. For example, a synapse delay may change only when a weight change occurs or if weights reach zero but not if they are at a maximum value. However, it may be advantageous to have independent functions so that these processes can be parallelized reducing the number and overlap of memory accesses.

DETERMINATION OF SYNAPTIC PLASTICITY

[0038] Neuroplasticity (or simply “plasticity”) is the capacity of neurons and neural networks in the brain to change their synaptic connections and behavior in response to new information, sensory stimulation, development, damage, or dysfunction. Plasticity is important to learning and memory in biology, as well as for computational neuroscience and neural networks. Various forms of plasticity have been studied, such as synaptic plasticity (e.g., according to the Hebbian theory), spike-timing-dependent plasticity (STDP), non-synaptic plasticity, activity-dependent plasticity, structural plasticity and homeostatic plasticity.

[0039] STDP is a learning process that adjusts the strength of synaptic connections between neurons. The connection strengths are adjusted based on the relative timing of a particular neuron's output and received input spikes (i.e., action potentials). Under the STDP process, long-term potentiation (LTP) may occur if an input spike to a certain neuron tends, on average, to occur immediately before that neuron's output spike. Then, that particular input is made somewhat stronger. On the other hand, long-term

depression (LTD) may occur if an input spike tends, on average, to occur immediately after an output spike. Then, that particular input is made somewhat weaker, and hence the name "spike-timing-dependent plasticity." Consequently, inputs that might be the cause of the postsynaptic neuron's excitation are made even more likely to contribute in the future, whereas inputs that are not the cause of the postsynaptic spike are made less likely to contribute in the future. The process continues until a subset of the initial set of connections remains, while the influence of all others is reduced to an insignificant level.

[0040] Because a neuron generally produces an output spike when many of its inputs occur within a brief period (i.e., being cumulative sufficient to cause the output), the subset of inputs that typically remains includes those that tended to be correlated in time. In addition, because the inputs that occur before the output spike are strengthened, the inputs that provide the earliest sufficiently cumulative indication of correlation will eventually become the final input to the neuron.

[0041] The STDP learning rule may effectively adapt a synaptic weight of a synapse connecting a presynaptic neuron to a postsynaptic neuron as a function of time difference between spike time t_{pre} of the presynaptic neuron and spike time t_{post} of the postsynaptic neuron (i.e., $t = t_{post} - t_{pre}$). A typical formulation of the STDP is to increase the synaptic weight (i.e., potentiate the synapse) if the time difference is positive (the presynaptic neuron fires before the postsynaptic neuron), and decrease the synaptic weight (i.e., depress the synapse) if the time difference is negative (the postsynaptic neuron fires before the presynaptic neuron).

[0042] In the STDP process, a change of the synaptic weight over time may be typically achieved using an exponential decay, as given by:

$$\Delta w(t) = \begin{cases} a_+ e^{-t/k_+} + \mu, & t > 0 \\ a_- e^{t/k_-}, & t < 0 \end{cases}, \quad (1)$$

where k_+ and k_- are time constants for positive and negative time difference, respectively, a_+ and a_- are corresponding scaling magnitudes, and μ is an offset that may be applied to the positive time difference and/or the negative time difference.

[0043] FIGURE 3 illustrates an exemplary diagram 300 of a synaptic weight change as a function of relative timing of presynaptic and postsynaptic spikes in accordance with the STDP. If a presynaptic neuron fires before a postsynaptic neuron, then a corresponding synaptic weight may be increased, as illustrated in a portion 302 of the graph 300. This weight increase can be referred to as an LTP of the synapse. It can be observed from the graph portion 302 that the amount of LTP may decrease roughly exponentially as a function of the difference between presynaptic and postsynaptic spike times. The reverse order of firing may reduce the synaptic weight, as illustrated in a portion 304 of the graph 300, causing an LTD of the synapse.

[0044] As illustrated in the graph 300 in FIGURE 3, a negative offset μ may be applied to the LTP (causal) portion 302 of the STDP graph. A point of cross-over 306 of the x-axis ($y=0$) may be configured to coincide with the maximum time lag for considering correlation for causal inputs from layer $i-1$. In the case of a frame-based input (i.e., an input that is in the form of a frame of a particular duration comprising spikes or pulses), the offset value μ can be computed to reflect the frame boundary. A first input spike (pulse) in the frame may be considered to decay over time either as modeled by a postsynaptic potential directly or in terms of the effect on neural state. If a second input spike (pulse) in the frame is considered correlated or relevant to a particular time frame, then the relevant times before and after the frame may be separated at that time frame boundary and treated differently in plasticity terms by offsetting one or more parts of the STDP curve such that the value in the relevant times may be different (e.g., negative for greater than one frame and positive for less than one frame). For example, the negative offset μ may be set to offset LTP such that the curve actually goes below zero at a pre-post time greater than the frame time and it is thus part of LTD instead of LTP.

NEURON MODELS AND OPERATION

[0045] There are some general principles for designing a useful spiking neuron model. A good neuron model may have rich potential behavior in terms of two computational regimes: coincidence detection and functional computation. Moreover, a good neuron model should have two elements to allow temporal coding: arrival time of inputs affects output time and coincidence detection can have a narrow time window. Finally, to be computationally attractive, a good neuron model may have a closed-form

solution in continuous time and stable behavior including near attractors and saddle points. In other words, a useful neuron model is one that is practical and that can be used to model rich, realistic and biologically-consistent behaviors, as well as be used to both engineer and reverse engineer neural circuits.

[0046] A neuron model may depend on events, such as an input arrival, output spike or other event whether internal or external. To achieve a rich behavioral repertoire, a state machine that can exhibit complex behaviors may be desired. If the occurrence of an event itself, separate from the input contribution (if any), can influence the state machine and constrain dynamics subsequent to the event, then the future state of the system is not only a function of a state and input, but rather a function of a state, event, and input.

[0047] In an aspect, a neuron n may be modeled as a spiking leaky-integrate-and-fire neuron with a membrane voltage $v_n(t)$ governed by the following dynamics:

$$\frac{dv_n(t)}{dt} = \alpha v_n(t) + \beta \sum_m w_{m,n} y_m(t - \Delta t_{m,n}), \quad (2)$$

where α and β are parameters, $w_{m,n}$ is a synaptic weight for the synapse connecting a presynaptic neuron m to a postsynaptic neuron n , and $y_m(t)$ is the spiking output of the neuron m that may be delayed by dendritic or axonal delay according to $\Delta t_{m,n}$ until arrival at the neuron n 's soma.

[0048] It should be noted that there is a delay from the time when sufficient input to a postsynaptic neuron is established until the time when the postsynaptic neuron actually fires. In a dynamic spiking neuron model, such as Izhikevich's simple model, a time delay may be incurred if there is a difference between a depolarization threshold v_t and a peak spike voltage v_{peak} . For example, in the simple model, neuron soma dynamics can be governed by the pair of differential equations for voltage and recovery, i.e.:

$$\frac{dv}{dt} = (k(v - v_t)(v - v_r) - u + I)/C, \quad (3)$$

$$\frac{du}{dt} = a(b(v - v_r) - u). \quad (4)$$

where v is a membrane potential, u is a membrane recovery variable, k is a parameter that describes time scale of the membrane potential v , a is a parameter that describes time scale of the recovery variable u , b is a parameter that describes sensitivity of the recovery variable u to the sub-threshold fluctuations of the membrane potential v , v_r is a membrane resting potential, I is a synaptic current, and C is a membrane's capacitance. In accordance with this model, the neuron is defined to spike when $v > v_{peak}$.

Hunzinger Cold Model

[0049] The Hunzinger Cold neuron model is a minimal dual-regime spiking linear dynamical model that can reproduce a rich variety of neural behaviors. The model's one- or two-dimensional linear dynamics can have two regimes, wherein the time constant (and coupling) can depend on the regime. In the sub-threshold regime, the time constant, negative by convention, represents leaky channel dynamics generally acting to return a cell to rest in a biologically-consistent linear fashion. The time constant in the supra-threshold regime, positive by convention, reflects anti-leaky channel dynamics generally driving a cell to spike while incurring latency in spike-generation.

[0050] As illustrated in FIGURE 4, the dynamics of the model 400 may be divided into two (or more) regimes. These regimes may be called the negative regime 402 (also interchangeably referred to as the leaky-integrate-and-fire (LIF) regime, not to be confused with the LIF neuron model) and the positive regime 404 (also interchangeably referred to as the anti-leaky-integrate-and-fire (ALIF) regime, not to be confused with the ALIF neuron model). In the negative regime 402, the state tends toward rest (v_-) at the time of a future event. In this negative regime, the model generally exhibits temporal input detection properties and other sub-threshold behavior. In the positive regime 404, the state tends toward a spiking event (v_s). In this positive regime, the model exhibits computational properties, such as incurring a latency to spike depending on subsequent input events. Formulation of dynamics in terms of events and separation of the dynamics into these two regimes are fundamental characteristics of the model.

[0051] Linear dual-regime bi-dimensional dynamics (for states v and u) may be defined by convention as:

$$\tau_\rho \frac{dv}{dt} = v + q_\rho \quad (5)$$

$$-\tau_u \frac{du}{dt} = u + r \quad (6)$$

where q_ρ and r are the linear transformation variables for coupling.

[0052] The symbol ρ is used herein to denote the dynamics regime with the convention to replace the symbol ρ with the sign “-” or “+” for the negative and positive regimes, respectively, when discussing or expressing a relation for a specific regime.

[0053] The model state is defined by a membrane potential (voltage) v and recovery current u . In basic form, the regime is essentially determined by the model state. There are subtle, but important aspects of the precise and general definition, but for the moment, consider the model to be in the positive regime 404 if the voltage v is above a threshold (v_+) and otherwise in the negative regime 402.

[0054] The regime-dependent time constants include τ_- which is the negative regime time constant, and τ_+ which is the positive regime time constant. The recovery current time constant τ_u is typically independent of regime. For convenience, the negative regime time constant τ_- is typically specified as a negative quantity to reflect decay so that the same expression for voltage evolution may be used as for the positive regime in which the exponent and τ_+ will generally be positive, as will be τ_u .

[0055] The dynamics of the two state elements may be coupled at events by transformations offsetting the states from their null-clines, where the transformation variables are:

$$q_\rho = -\tau_\rho \beta u - v_\rho \quad (7)$$

$$r = \delta(v + \varepsilon) \quad (8)$$

where δ , ε , β and v_- , v_+ are parameters. The two values for v_ρ are the base for reference voltages for the two regimes. The parameter v_- is the base voltage for the negative regime, and the membrane potential will generally decay toward v_- in the negative regime. The parameter v_+ is the base voltage for the positive regime, and the membrane potential will generally tend away from v_+ in the positive regime.

[0056] The null-clines for v and u are given by the negative of the transformation variables q_ρ and r , respectively. The parameter δ is a scale factor controlling the slope of the u null-cline. The parameter ε is typically set equal to $-v_-$. The parameter β is a resistance value controlling the slope of the v null-clines in both regimes. The τ_ρ time-constant parameters control not only the exponential decays, but also the null-cline slopes in each regime separately.

[0057] The model may be defined to spike when the voltage v reaches a value v_s . Subsequently, the state may be reset at a reset event (which may be one and the same as the spike event):

$$v = \hat{v}_- \quad (9)$$

$$u = u + \Delta u \quad (10)$$

where \hat{v}_- and Δu are parameters. The reset voltage \hat{v}_- is typically set to v_- .

[0058] By a principle of momentary coupling, a closed form solution is possible not only for state (and with a single exponential term), but also for the time to reach a particular state. The close form state solutions are:

$$v(t + \Delta t) = (v(t) + q_\rho) e^{\frac{\Delta t}{\tau_\rho}} - q_\rho \quad (11)$$

$$u(t + \Delta t) = (u(t) + r) e^{-\frac{\Delta t}{\tau_u}} - r \quad (12)$$

[0059] Therefore, the model state may be updated only upon events, such as an input (presynaptic spike) or output (postsynaptic spike). Operations may also be performed at any particular time (whether or not there is input or output).

[0060] Moreover, by the momentary coupling principle, the time of a postsynaptic spike may be anticipated so the time to reach a particular state may be determined in advance without iterative techniques or Numerical Methods (e.g., the Euler numerical method). Given a prior voltage state v_0 , the time delay until voltage state v_f is reached is given by:

$$\Delta t = \tau_\rho \log \frac{v_f + q_\rho}{v_0 + q_\rho} \quad (13)$$

[0061] If a spike is defined as occurring at the time the voltage state v reaches v_s , then the closed-form solution for the amount of time, or relative delay, until a spike occurs as measured from the time that the voltage is at a given state v is:

$$\Delta t_s = \begin{cases} \tau_+ \log \frac{v_s + q_+}{v + q_+} & \text{if } v > \hat{v}_+ \\ \infty & \text{otherwise} \end{cases} \quad (14)$$

where \hat{v}_+ is typically set to parameter v_+ , although other variations may be possible.

[0062] The above definitions of the model dynamics depend on whether the model is in the positive or negative regime. As mentioned, the coupling and the regime ρ may be computed upon events. For purposes of state propagation, the regime and coupling (transformation) variables may be defined based on the state at the time of the last (prior) event. For purposes of subsequently anticipating spike output time, the regime and coupling variable may be defined based on the state at the time of the next (current) event.

[0063] There are several possible implementations of the Cold model, and executing the simulation, emulation or model in time. This includes, for example, event-update, step-event update, and step-update modes. An event update is an update where states are updated based on events or “event update” (at particular moments). A step update is an update when the model is updated at intervals (e.g., 1ms). This does not necessarily utilize iterative methods or Numerical methods. An event-based implementation is also possible at a limited time resolution in a step-based simulator by only updating the model if an event occurs at or between steps or by “step-event” update.

Auditory Source Separation in a Spiking Neural Network

[0064] The separation of mixed signals according to auditory sources may be beneficial and broadly applicable. For example, auditory source segregation may provide speech segregation, aid in speech enhancement and aid in speech activity detection. However, monaural, unsupervised, online auditory source segregation is a challenge for computer audition. Accordingly, aspects of the present disclosure are directed to spike based implementations of temporal coherence based source segregation.

[0065] FIGURE 5 illustrates an example implementation 500 of the aforementioned auditory source segregation using a general-purpose processor 502 in accordance with certain aspects of the present disclosure. Variables (neural signals), synaptic weights, system parameters associated with a computational network (neural network), delays, and frequency bin information may be stored in a memory block 504, while instructions executed at the general-purpose processor 502 may be loaded from a program memory 506. In an aspect of the present disclosure, the instructions loaded into the general-purpose processor 502 may comprise code for selecting an audio attribute of an audio signal, representing a portion of the audio attribute that is dominated by a single source as a spiking event, representing a remaining portion of the audio signal as a spiking event, and/or determining whether the remaining portion coincides with the single source based on a coincidence of the spiking events.

[0066] FIGURE 6 illustrates an example implementation 600 of the aforementioned auditory source segregation where a memory 602 can be interfaced via an interconnection network 604 with individual (distributed) processing units (neural processors) 606 of a computational network (neural network) in accordance with certain aspects of the present disclosure. Variables (neural signals), synaptic weights, system parameters associated with the computational network (neural network) delays, frequency bin information, audio signals, and/or audio attribute information may be stored in the memory 602, and may be loaded from the memory 602 via connection(s) of the interconnection network 604 into each processing unit (neural processor) 606. In an aspect of the present disclosure, the processing unit 606 may be configured to select an audio attribute of an audio signal, represent a portion of the audio attribute that is dominated by a single source as a spiking event, represent a remaining portion of the

audio signal as a spiking event, and/or determine whether the remaining portion coincides with the single source based on a coincidence of the spiking events.

[0067] FIGURE 7 illustrates an example implementation 700 of the aforementioned auditory source segregation. As illustrated in FIGURE 7, one memory bank 702 may be directly interfaced with one processing unit 704 of a computational network (neural network). Each memory bank 702 may store variables (neural signals), synaptic weights, and/or system parameters associated with a corresponding processing unit (neural processor) 704 delays, frequency bin information, audio signals and/or audio attribute information. In an aspect of the present disclosure, the processing unit 704 may be configured to select an audio attribute of an audio signal, represent a portion of the audio attribute that is dominated by a single source as a spiking event, represent a remaining portion of the audio signal as a spiking event, and/or determine whether the remaining portion coincides with the single source based on a coincidence of the spiking events.

[0068] FIGURE 8 illustrates an example implementation of a neural network 800 in accordance with certain aspects of the present disclosure. As illustrated in FIGURE 8, the neural network 800 may have multiple local processing units 802 that may perform various operations of methods described herein. Each local processing unit 802 may comprise a local state memory 804 and a local parameter memory 806 that store parameters of the neural network. In addition, the local processing unit 802 may have a local (neuron) model program (LMP) memory 808 for storing a local model program, a local learning program (LLP) memory 810 for storing a local learning program, and a local connection memory 812. Furthermore, as illustrated in FIGURE 8, each local processing unit 802 may be interfaced with a configuration processor unit 814 for providing configurations for local memories of the local processing unit, and with a routing unit 816 that provide routing between the local processing units 802.

[0069] In one configuration, a neuron model is configured for selecting an audio attribute of an audio signal, representing a portion of the audio attribute that is dominated by a single source as a spiking event, representing a remaining portion of the audio signal as a spiking event, and/or determining whether the remaining portion coincides with the single source based on a coincidence of the spiking events. The neuron model includes a selecting means, means for representing a portion of the audio

attribute that is dominated by a single source as a spiking event, means for representing a remaining portion of the audio signal as a spiking event and determining means. In one aspect, the selecting means, means for representing a portion of the audio attribute, means for representing a remaining portion of the audio signal, and/or determining means may be the general-purpose processor 502, program memory 506, memory block 504, memory 602, interconnection network 604, processing units 606, processing unit 704, local processing units 802, and/or the routing connection processing elements 816 configured to perform the functions recited. In another configuration, the aforementioned means may be any module or any apparatus configured to perform the functions recited by the aforementioned means.

[0070] According to certain aspects of the present disclosure, each local processing unit 802 may be configured to determine parameters of the neural network based upon desired one or more functional features of the neural network, and develop the one or more functional features towards the desired functional features as the determined parameters are further adapted, tuned and updated.

[0071] FIGURE 9 is a block diagram illustrating an exemplary architecture 900 for auditory source segregation in accordance with aspects of the present disclosure. The architecture includes an input 902 to supply an audio signal. The audio signal may be analog or digital. In some aspects, the audio signal may be provided via an audio input device such as a microphone, a storage device such as digital audio tape (DAT), compact disc (CD), Digital video disc (DVD), Blu-ray Disc (BD) or via streaming media, for example.

[0072] In some aspects, the audio signal may comprise a cortical feature mixture. The cortical feature mixture may include different output channels associated with different output frequencies. For example, the audio signal may include a mixture of voices from one or more speakers, and/or other sounds (e.g., background noise, music or animals sounds) of an auditory scene (e.g., a crowded bus or sporting event).

[0073] The audio signal may be supplied to an audio processor 904. In some aspects, the audio processor may process the audio signal to determine an audio attribute that is dominated by a single or target source. An audio attribute may for example, comprise pitch, timbre, harmonics, rhythm, loudness, speed (rate of speaking),

spatial location or the like. In one example, the audio processor 904 may comprise a bank of band pass filters with center frequencies spanning a predetermined range. The audio processor 904 may identify a pitch (i.e., an audio attribute) that is dominated by a target source (e.g., voice of a male or female speaker). In this example, a dominated audio attribute or anchor may be identified as the pitch having the highest frequency in a given range (e.g., with center frequencies between 180-200 Hz for a point in time (e.g., every 25 msec)). The anchor may provide an indication of a channel of the audio signal for a particular source. A pitch track may be produced based on a collection of raw power values corresponding to the indicated channel of the audio signal.

[0074] In some aspects, the audio signal and/or a portion of the audio signal corresponding to the dominated audio attribute (e.g., pitch track) may be subject to rectification and normalization processing. For example, the dominated audio attribute portion may be normalized by dividing the raw power values by a predetermined raw power value (e.g., a maximum power value) for an entire interval (e.g., given range). In another example, the dominated audio attribute portion may be rectified by setting all negative values to 0.

[0075] The dominated audio attribute portion (e.g., pitch track) may be supplied to a layer of neurons 906 along with the audio signal. In one exemplary aspect, the layer of neurons 906 may comprise leaky integrate and fire (LIF) neurons. The neurons 906 may be configured to represent the dominated audio attribute portion and the audio signal as spikes or spike events. For example, the LIF neurons may be configured to spike once for each oscillation. In a further example, a time constant (e.g., τ_+) for the LIF neurons may be set such that a positive oscillation of the input will generate one or more spikes at the peak of the oscillation.

[0076] The output spikes representing the dominated audio attribute portion (e.g., pitch track) and the output spikes representing the audio signal may be supplied as inputs to a layer of Coincidence Detector Object (CDO) neurons 908. The layer of CDO neurons 908 may be configured to compare the timing of the spikes representing the dominated audio attribute and timing of the spikes representing the audio signal. When the spikes representing the dominated audio attribute and the spikes representing the audio signal are coincident or match, the CDO neurons 908 may spike. In some aspects, the CDO neurons 908 may also spike when the timing difference between the

spikes representing the dominated audio attribute and the spikes representing the audio signal is below a predefined threshold (e.g., within a timing or temporal window). The output spikes of the CDO neurons 908 may represent a portion of the audio signal, which is attributable to the target source. In this way, the CDO neurons 908 may be used to search for and segregate a portion of the audio signal attributable to the target source.

[0077] Accordingly, a portion of the audio signal attributed to a target source may be extracted or segregated from an audio signal or auditory scene with multiple sound sources on the fly without prior knowledge of the target source. This may also be accomplished without the benefit of training on noise background.

[0078] FIGURE 10 is a block diagram 1000 illustrating an exemplary architecture for auditory source segregation in accordance with aspects of the present disclosure. The architecture may be configured to segregate audio from multiple sources included in an audio signal. Referring to FIGURE 10, the architecture may include an input device 1006 for providing an audio signal. The architecture may also include an input device 1002 for providing an audio attribute associated with a sound source A and an input device 1004 for providing an audio attribute associated with a sound source B. Although, three separate inputs are shown, this is merely for ease of explanation, and any number of input devices could be used. For example, a single input could also be used to provide each of the input signals.

[0079] In some aspects, the audio attribute of speaker A and the audio attribute of sound source B may be determined on the fly without prior knowledge of the speakers. Of course, this is merely exemplary, and in some aspects, the audio attributes for sound source A and sound source B may be determined in advance and stored in a storage device. In addition, the audio attribute for sound source A and the audio attribute for sound source B may be different. For example, the audio attribute for sound source A may be speech location, while the audio attribute for sound source B may be loudness.

[0080] As shown in FIGURE 10, an audio signal may be provided via the input device 1006. The audio signal may comprise a cortical feature mixture. The cortical feature mixture may include different output channels associated with different output frequencies. For example, in some aspects, the audio signal may include a mixture of

voices from one or more speakers, and/or other sounds (e.g., background noise, music, or animal sounds) of an auditory scene. The auditory scene may comprise a mixture of sounds of a crowded bus or sporting event, for instance. In an exemplary configuration, the audio source is processes through a band of band pass filters with center frequencies spanning a range. The output may be collected into a spectrogram, which may be filtered along the frequency axis for multiple scales to generate a multi-resolution representation.

[0081] An audio attribute (e.g., pitch) for sound source A and sound source B are supplied via input devices 1002 and 1004. In some aspects, an estimate of the audio attribute for the sound sources may be used. For example, the template technique for pitch extraction can be applied.

[0082] In one configuration, the audio attribute is identified as band pass filter output that is highest in a specified range (e.g., with center frequencies between 180-200 Hz) for a time period (e.g., 25 ms). For each identified pitch channel, raw power values can be collected from the corresponding channel in a spectrogram and at the corresponding time, yielding a one dimensional array. For time points in which there is no identified pitch channel, the array carries a value of 0. The pitch track may be further processed with a band pass filter at a new rate (e.g., 2 Hz). The values may be normalized by dividing all values by the predetermined value (e.g., maximum value) of the entire speech interval. Moreover, the output may be rectified, i.e., all negative values are set to 0. The filtering may be repeated at other frequencies (e.g., 4 Hz, 8 Hz 16 Hz). The output may be collected into a spectrogram, which may be filtered along the frequency axis for multiple scales. Although, FIGURE 10 shows an audio attribute for two sound sources, this is merely exemplary and not limiting. Rather, any number of audio attributes for any number of sound sources may be utilized.

[0083] The audio attribute may comprise a channel of the audio signal. In some aspects, the audio attribute may comprise pitch, timbre, harmonics, rhythm, loudness, speed (rate of speaking), speech location or the like. In some aspects, the audio attributes for the sound sources and/or the audio signal may be rectified and normalized.

[0084] The audio attributes and the audio signal may be supplied to a layer of neurons 1008. The layer of neurons 1008 may be configured to represent each of the

audio attributes and the audio signal as spikes. In some aspects, the layer of neurons may comprise leaky integrate and fire (LIF) neurons, for example.

[0085] The spike representations of the audio attributes may be respectively supplied to a corresponding layer of CDO neurons 1010, 1012 as a first input. The spike representation of the audio signal may be supplied as a second input to each of the layers of CDO neurons 1010, 1012. The neuron model may have a time constant that is set such that a positive oscillation of the input will generate one or more spikes at the peak of oscillation. The oscillation occurs at the frequency of the second filter (e.g., 2 Hz). The process may repeat for each of the other filter frequencies (e.g., 4 Hz, 8 Hz, 16 Hz).

[0086] The CDO neurons (1010, 1012) may, in turn determine whether there is a correlation between the audio signal and the audio attribute for sound source A and/or the audio attribute for sound source B. That is, the CDO neurons 1010, 1012 may be configured to determine whether the spikes representing the audio attributes are coincident with the spikes that represent the audio signal. When the spikes are coincident or within a predetermined range of each other, the CDO neurons may output a spike. Accordingly, the CDO neurons may provide an indication of whether the audio attribute for a sound source is present within the audio signal. Further, when the audio attribute for a sound source is present within the audio signal, a portion of the audio signal attributed to that sound source may be extracted or segregated from the audio signal.

[0087] Although, FIGURE 10 shows architecture which includes an audio attribute (e.g., pitch) for two sound sources and two CDO neurons, this is merely exemplary for ease of illustration and explanation and the architecture may be configured to separate any number of sources from a mixture. That is, the architecture may be configured to separate one or more sources from an audio signal including more than one source. For example, the architecture may be configured to segregate one or more voices from a noisy audio signal supplied from a crowded restaurant or bus.

[0088] Furthermore, while the present disclosure has highlighted aspects related to audio source segregation, other applications may be realized using the architecture and features disclosed herein. These applications include, but are not limited to, speech

activity detection, speech recognition, speech coding, and audio enhancements. For example, in some aspects, the segregated audio signals for sound source A and sound source B shown in FIGURE 10 may be recombined and decoded to produce a composite audio stream without other sounds included in the input audio signal such as background noise. In this way, speech enhancements such as noise reduction may be realized.

[0089] FIGURE 11 illustrates an exemplary method 1100 for auditory source segregation in accordance with aspects of the present disclosure. In block 1102, the neuron model selects an audio attribute of an audio signal. In some aspects, the audio attribute may be a pitch, volume, spatial location, timbre, tone, rate of speaking or any other audio attribute.

[0090] In block 1104, the neuron model represents a portion of the audio attribute that is dominated by a single source as a spiking event. In some aspects, an audio attribute dominated by multiple sources may also be represented as a spike event.

[0091] In block 1106, the neuron model represents a remaining portion of the audio signal as spiking events. Furthermore, in block 1108, the neuron model determines whether the remaining portion coincides with the single source based on a coincidence of the spiking events.

[0092] In some aspects, the single source may be an anchor source. The anchor may be any audio feature or channel dominated by a source. The anchor may be processed using band pass filters to identify a filter output within a specified range. Further, the processing may be conducted on a periodic basis.

[0093] The various operations of methods described above may be performed by any suitable means capable of performing the corresponding functions. The means may include various hardware and/or software component(s) and/or module(s), including, but not limited to, a circuit, an application specific integrated circuit (ASIC), or processor. Generally, where there are operations illustrated in the figures, those operations may have corresponding counterpart means-plus-function components with similar numbering.

[0094] As used herein, the term “determining” encompasses a wide variety of actions. For example, “determining” may include calculating, computing, processing, deriving, investigating, looking up (e.g., looking up in a table, a database or another data structure), ascertaining and the like. Additionally, “determining” may include receiving (e.g., receiving information), accessing (e.g., accessing data in a memory) and the like. Furthermore, “determining” may include resolving, selecting, choosing, establishing and the like.

[0095] As used herein, a phrase referring to “at least one of” a list of items refers to any combination of those items, including single members. As an example, “at least one of: a, b, or c” is intended to cover: a, b, c, a-b, a-c, b-c, and a-b-c.

[0096] The various illustrative logical blocks, modules and circuits described in connection with the present disclosure may be implemented or performed with a general purpose processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array signal (FPGA) or other programmable logic device (PLD), discrete gate or transistor logic, discrete hardware components or any combination thereof designed to perform the functions described herein. A general-purpose processor may be a microprocessor, but in the alternative, the processor may be any commercially available processor, controller, microcontroller or state machine. A processor may also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration.

[0097] The steps of a method or algorithm described in connection with the present disclosure may be embodied directly in hardware, in a software module executed by a processor, or in a combination of the two. A software module may reside in any form of storage medium that is known in the art. Some examples of storage media that may be used include random access memory (RAM), read only memory (ROM), flash memory, erasable programmable read-only memory (EPROM), electrically erasable programmable read-only memory (EEPROM), registers, a hard disk, a removable disk, a CD-ROM and so forth. A software module may comprise a single instruction, or many instructions, and may be distributed over several different code segments, among different programs, and across multiple storage media. A storage medium may be coupled to a processor such that the processor can read information from, and write

information to, the storage medium. In the alternative, the storage medium may be integral to the processor.

[0098] The methods disclosed herein comprise one or more steps or actions for achieving the described method. The method steps and/or actions may be interchanged with one another without departing from the scope of the claims. In other words, unless a specific order of steps or actions is specified, the order and/or use of specific steps and/or actions may be modified without departing from the scope of the claims.

[0099] The functions described herein may be implemented in hardware, software, firmware, or any combination thereof. If implemented in hardware, an example hardware configuration may comprise a processing system in a device. The processing system may be implemented with a bus architecture. The bus may include any number of interconnecting buses and bridges depending on the specific application of the processing system and the overall design constraints. The bus may link together various circuits including a processor, machine-readable media, and a bus interface. The bus interface may be used to connect a network adapter, among other things, to the processing system via the bus. The network adapter may be used to implement signal processing functions. For certain aspects, a user interface (e.g., keypad, display, mouse, joystick, etc.) may also be connected to the bus. The bus may also link various other circuits such as timing sources, peripherals, voltage regulators, power management circuits, and the like, which are well known in the art, and therefore, will not be described any further.

[00100] The processor may be responsible for managing the bus and general processing, including the execution of software stored on the machine-readable media. The processor may be implemented with one or more general-purpose and/or special-purpose processors. Examples include microprocessors, microcontrollers, DSP processors, and other circuitry that can execute software. Software shall be construed broadly to mean instructions, data, or any combination thereof, whether referred to as software, firmware, middleware, microcode, hardware description language, or otherwise. Machine-readable media may include, by way of example, random access memory (RAM), flash memory, read only memory (ROM), programmable read-only memory (PROM), erasable programmable read-only memory (EPROM), electrically erasable programmable Read-only memory (EEPROM), registers, magnetic disks,

optical disks, hard drives, or any other suitable storage medium, or any combination thereof. The machine-readable media may be embodied in a computer-program product. The computer-program product may comprise packaging materials.

[00101] In a hardware implementation, the machine-readable media may be part of the processing system separate from the processor. However, as those skilled in the art will readily appreciate, the machine-readable media, or any portion thereof, may be external to the processing system. By way of example, the machine-readable media may include a transmission line, a carrier wave modulated by data, and/or a computer product separate from the device, all which may be accessed by the processor through the bus interface. Alternatively, or in addition, the machine-readable media, or any portion thereof, may be integrated into the processor, such as the case may be with cache and/or general register files. Although the various components discussed may be described as having a specific location, such as a local component, they may also be configured in various ways, such as certain components being configured as part of a distributed computing system.

[00102] The processing system may be configured as a general-purpose processing system with one or more microprocessors providing the processor functionality and external memory providing at least a portion of the machine-readable media, all linked together with other supporting circuitry through an external bus architecture. Alternatively, the processing system may comprise one or more neuromorphic processors for implementing the neuron models and models of neural systems described herein. As another alternative, the processing system may be implemented with an application specific integrated circuit (ASIC) with the processor, the bus interface, the user interface, supporting circuitry, and at least a portion of the machine-readable media integrated into a single chip, or with one or more field programmable gate arrays (FPGAs), programmable logic devices (PLDs), controllers, state machines, gated logic, discrete hardware components, or any other suitable circuitry, or any combination of circuits that can perform the various functionality described throughout this disclosure. Those skilled in the art will recognize how best to implement the described functionality for the processing system depending on the particular application and the overall design constraints imposed on the overall system.

[00103] The machine-readable media may comprise a number of software modules. The software modules include instructions that, when executed by the processor, cause the processing system to perform various functions. The software modules may include a transmission module and a receiving module. Each software module may reside in a single storage device or be distributed across multiple storage devices. By way of example, a software module may be loaded into RAM from a hard drive when a triggering event occurs. During execution of the software module, the processor may load some of the instructions into cache to increase access speed. One or more cache lines may then be loaded into a general register file for execution by the processor. When referring to the functionality of a software module below, it will be understood that such functionality is implemented by the processor when executing instructions from that software module.

[00104] If implemented in software, the functions may be stored or transmitted over as one or more instructions or code on a computer-readable medium. Computer-readable media include both computer storage media and communication media including any medium that facilitates transfer of a computer program from one place to another. A storage medium may be any available medium that can be accessed by a computer. By way of example, and not limitation, such computer-readable media can comprise RAM, ROM, EEPROM, CD-ROM or other optical disk storage, magnetic disk storage or other magnetic storage devices, or any other medium that can be used to carry or store desired program code in the form of instructions or data structures and that can be accessed by a computer. Also, any connection is properly termed a computer-readable medium. For example, if the software is transmitted from a website, server, or other remote source using a coaxial cable, fiber optic cable, twisted pair, digital subscriber line (DSL), or wireless technologies such as infrared (IR), radio, and microwave, then the coaxial cable, fiber optic cable, twisted pair, DSL, or wireless technologies such as infrared, radio, and microwave are included in the definition of medium. Disk and disc, as used herein, include compact disc (CD), laser disc, optical disc, digital versatile disc (DVD), floppy disk, and Blu-ray® disc where disks usually reproduce data magnetically, while discs reproduce data optically with lasers. Thus, in some aspects computer-readable media may comprise non-transitory computer-readable media (e.g., tangible media). In addition, for other aspects computer-readable media

may comprise transitory computer- readable media (e.g., a signal). Combinations of the above should also be included within the scope of computer-readable media.

[00105] Thus, certain aspects may comprise a computer program product for performing the operations presented herein. For example, such a computer program product may comprise a computer-readable medium having instructions stored (and/or encoded) thereon, the instructions being executable by one or more processors to perform the operations described herein. For certain aspects, the computer program product may include packaging material.

[00106] Further, it should be appreciated that modules and/or other appropriate means for performing the methods and techniques described herein can be downloaded and/or otherwise obtained by a user terminal and/or base station as applicable. For example, such a device can be coupled to a server to facilitate the transfer of means for performing the methods described herein. Alternatively, various methods described herein can be provided via storage means (e.g., RAM, ROM, a physical storage medium such as a compact disc (CD) or floppy disk, etc.), such that a user terminal and/or base station can obtain the various methods upon coupling or providing the storage means to the device. Moreover, any other suitable technique for providing the methods and techniques described herein to a device can be utilized.

[00107] It is to be understood that the claims are not limited to the precise configuration and components illustrated above. Various modifications, changes and variations may be made in the arrangement, operation and details of the methods and apparatus described above without departing from the scope of the claims.

CLAIMS**WHAT IS CLAIMED IS:**

1. A method of audio source segregation, comprising:
 - selecting an audio attribute of an audio signal;
 - representing a portion of the audio attribute dominated by a single source as at least one source spiking event;
 - representing a remaining portion of the audio signal as at least one audio signal spiking event; and
 - determining whether the remaining portion coincides with the single source based at least in part on coincidence of the at least one source spiking event and the at least one audio signal spiking event.
2. The method of claim 1, in which the determining is based at least in part on the at least one source spiking event and the at least one audio signal spiking event occurring within a temporal window.
3. The method of claim 1, in which the audio attribute is pitch.
4. The method of claim 1, in which the audio attribute is spatial location.
5. The method of claim 1, in which the determining is based at least in part on a rate of speaking.
6. The method of claim 1, in which the single source comprises an anchor and the anchor is processed using at least one band pass filter to identify a band pass filter output within a specified range.
7. The method of claim 6, in which the processing is conducted on a periodic basis.
8. An apparatus for audio source segregation, comprising:
 - a memory; and
 - at least one processor coupled to the memory, the at least one processor being configured:

to select an audio attribute of an audio signal;
to represent a portion of the audio attribute that is dominated by a single source as at least one source spiking event;
to represent a remaining portion of the audio signal as at least one audio signal spiking event; and
to determine whether the remaining portion coincides with the single source based at least in part on coincidence of the at least one source spiking event and the at least one audio signal spiking event.

9. The apparatus of claim 8, in which the at least one processor is further configured to determine whether the remaining portion coincides with the single source based at least in part on the at least one source spiking event and the at least one audio signal spiking event occurring within a temporal window.

10. The apparatus of claim 8, in which the audio attribute is pitch.

11. The apparatus of claim 8, in which the audio attribute is spatial location.

12. The apparatus of claim 8, in which the at least one processor is further configured to determine whether the remaining portion coincides with the single source based at least in part on a rate of speaking.

13. The apparatus of claim 8, in which the single source comprises an anchor and the at least one processor is further configured to process the anchor using at least one band pass filter to identify a band pass filter output within a specified range.

14. The apparatus of claim 13, in which the at least one processor is further configured to process the anchor on a periodic basis.

15. An apparatus for audio source segregation, comprising:
means for selecting an audio attribute of an audio signal;
means for representing a portion of the audio attribute that is dominated by a single source as at least one source spiking event;

means for representing a remaining portion of the audio signal as at least one audio signal spiking event; and

means for determining whether the remaining portion coincides with the single source based at least in part on coincidence of the at least one source spiking event and the at least one audio signal spiking event.

16. The apparatus of claim 15, in which the determining means further determines whether the remaining portion coincides with the single source based at least in part on the at least one source spiking event and the at least one audio signal spiking event occurring within a temporal window.

17. The apparatus of claim 15, in which the audio attribute is pitch.

18. A computer program product for audio source segregation, comprising:
a non-transitory computer readable medium having encoded thereon program code, the program code comprising:

program code to select an audio attribute of an audio signal;

program code to represent a portion of the audio attribute that is dominated by a single source as at least one source spiking event;

program code to represent a remaining portion of the audio signal as at least one audio signal spiking event; and

program code to determine whether the remaining portion coincides with the single source based at least in part on coincidence of the at least one source spiking event and the at least one audio signal spiking event.

19. The computer program product of claim 18, further comprising program code to determine whether the remaining portion coincides with the single source based at least in part on the at least one source spiking event and the at least one audio signal spiking event occurring within a temporal window.

20. The computer program product of claim 18, in which the audio attribute is pitch.

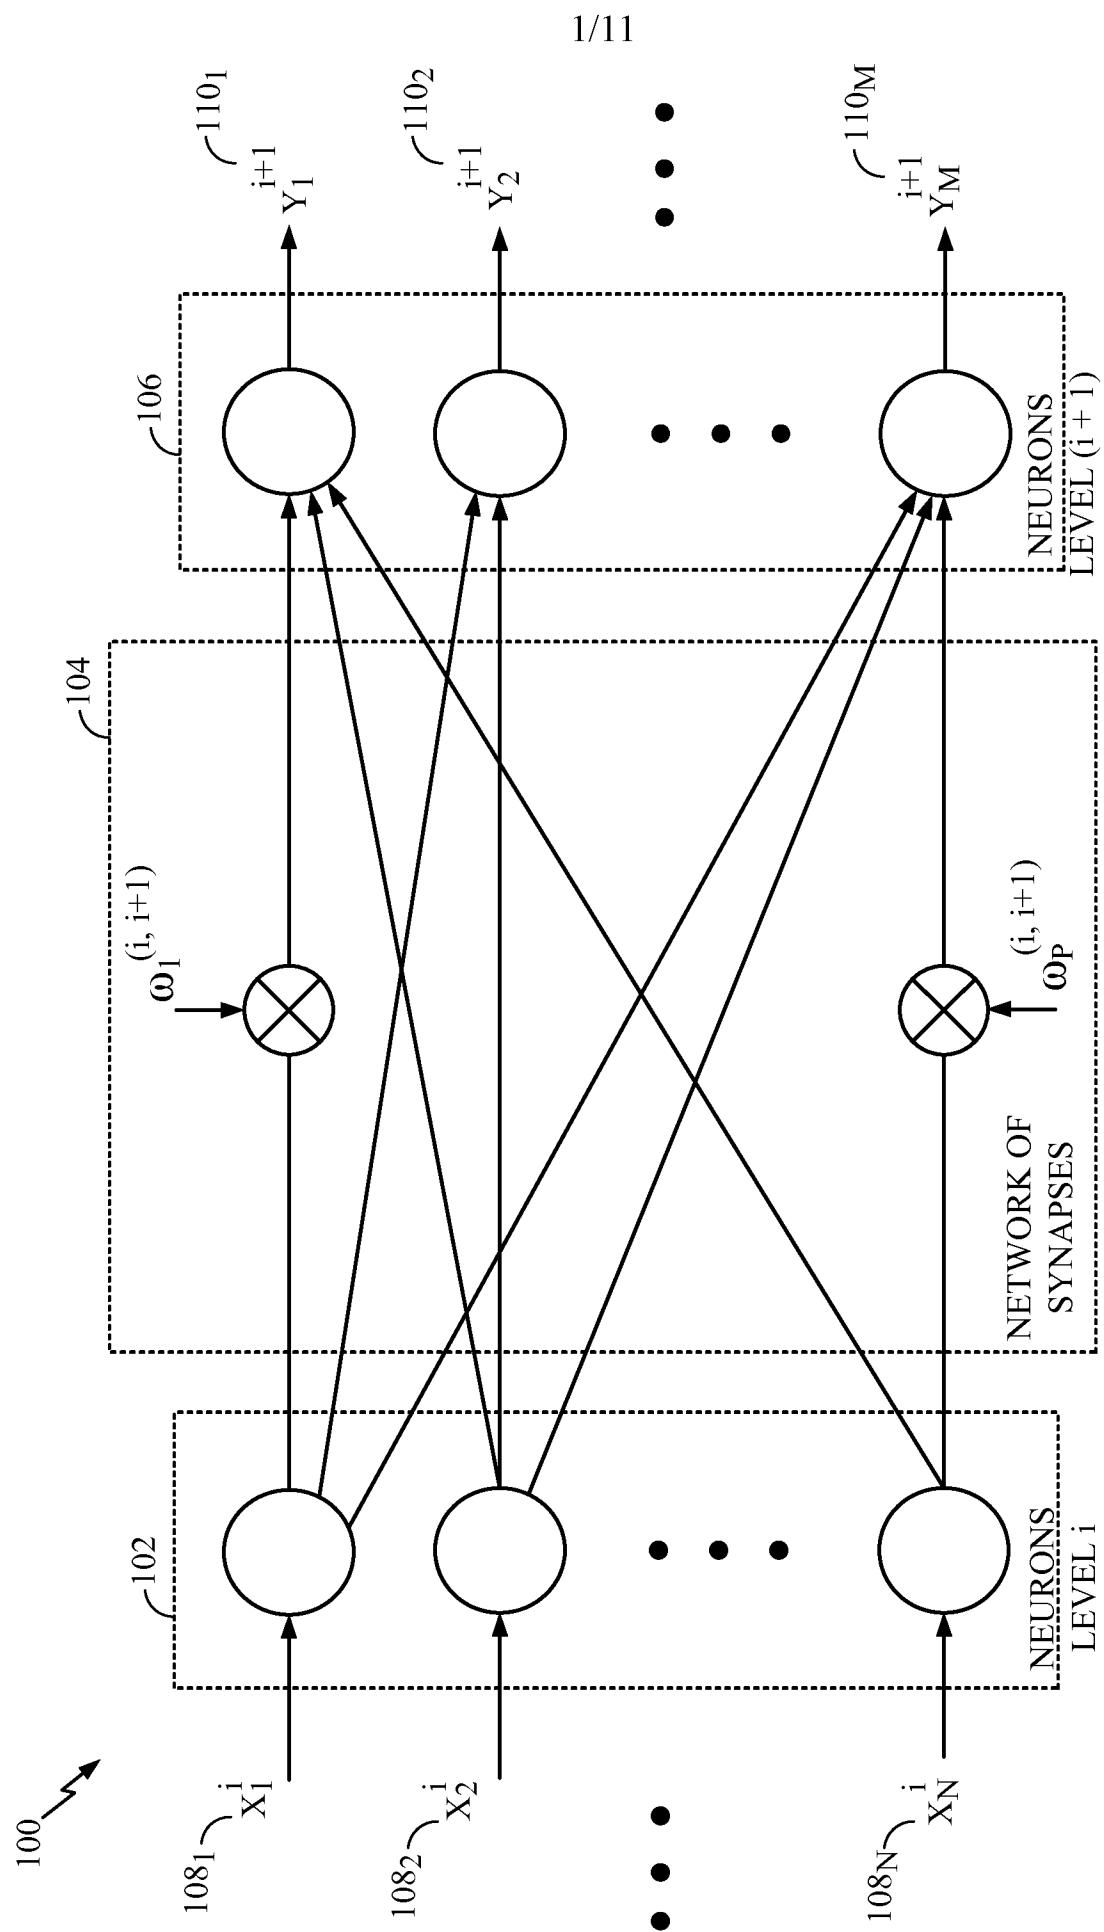
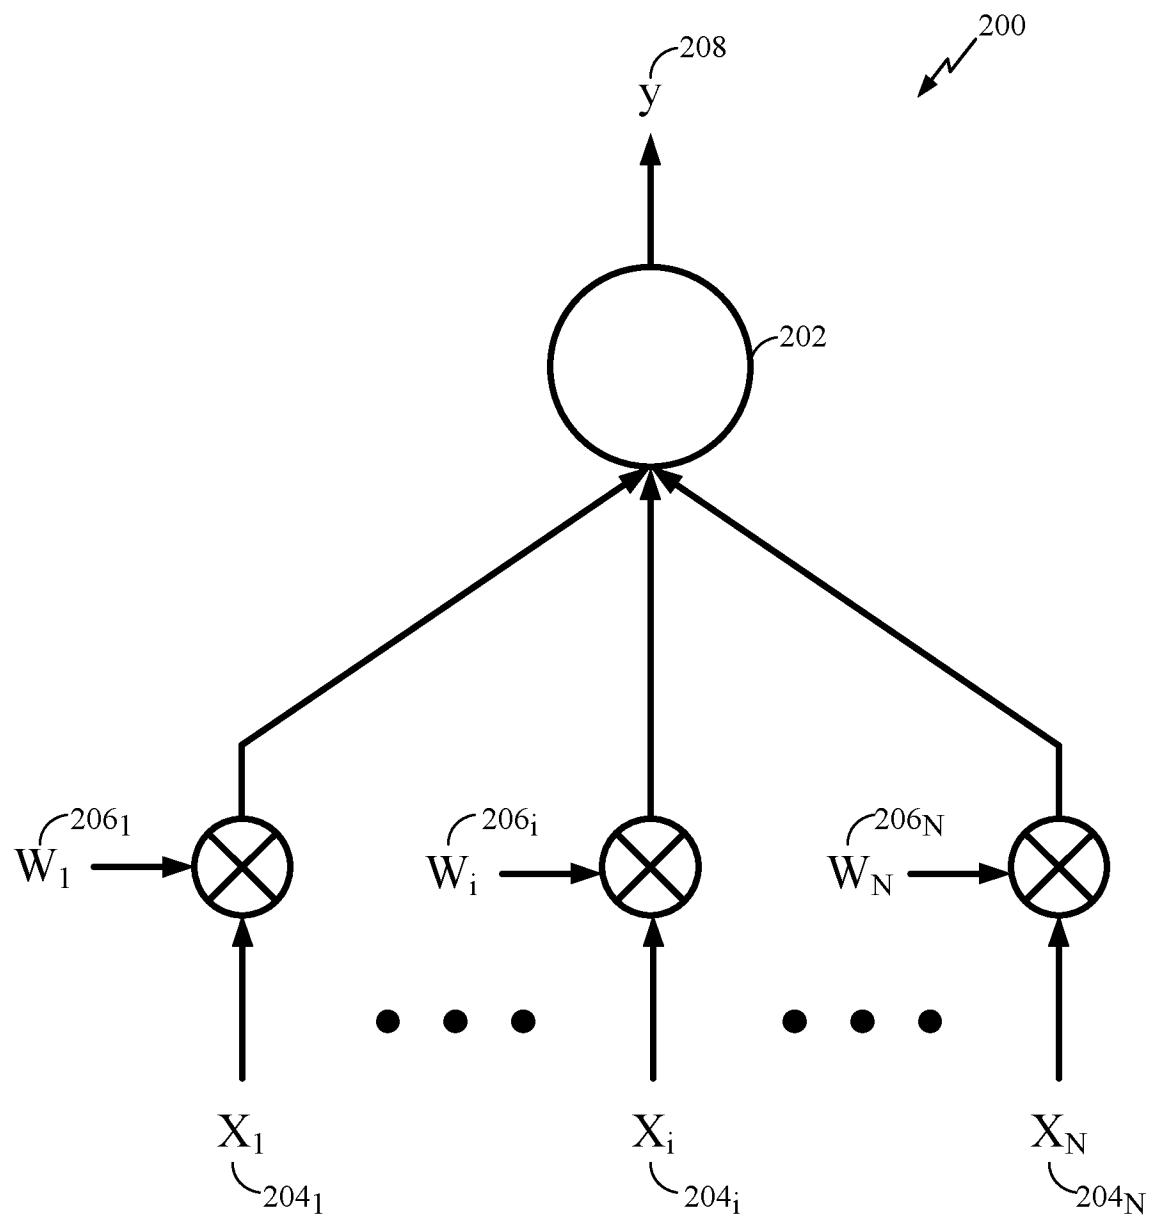
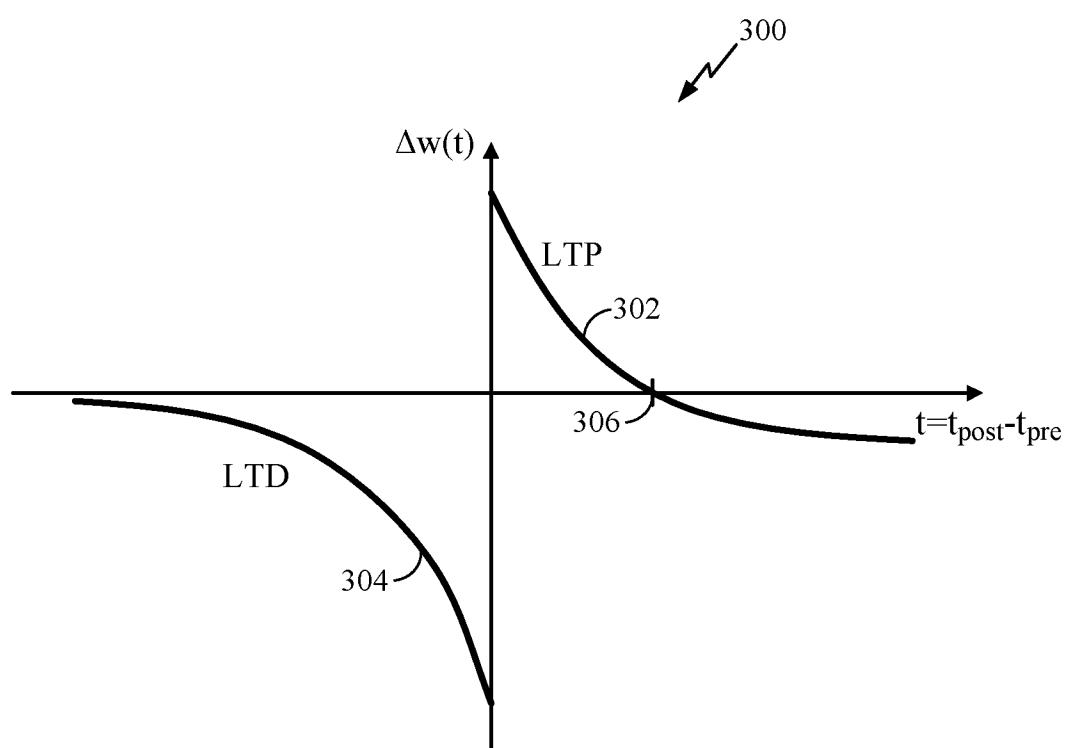


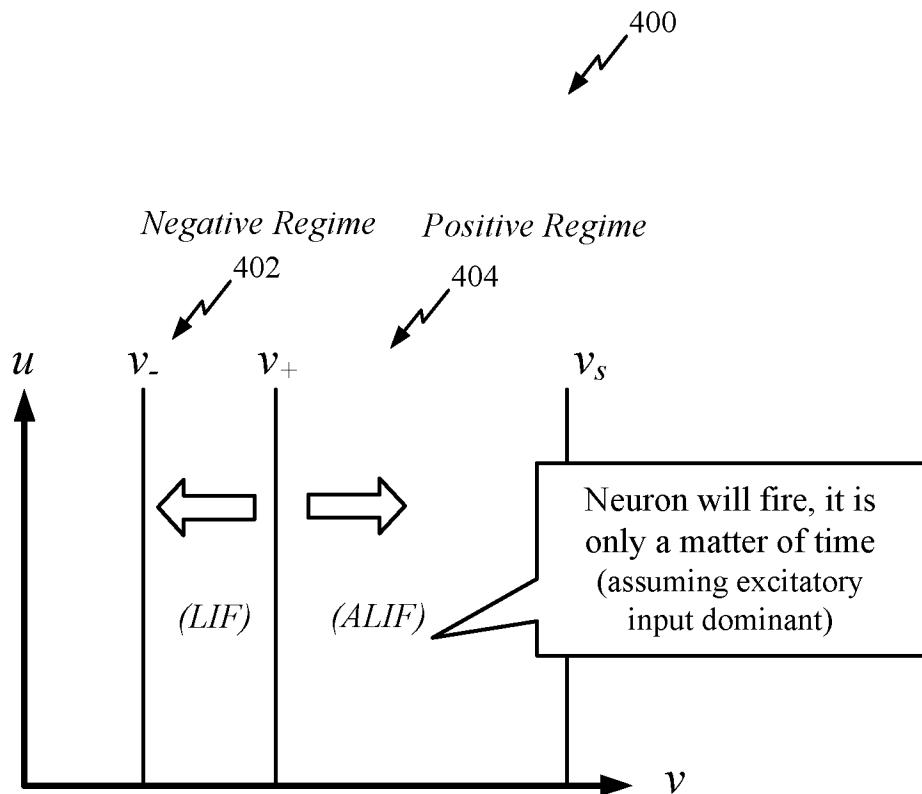
FIG. 1

**FIG. 2**

3/11

**FIG. 3**

4/11

**FIG. 4**

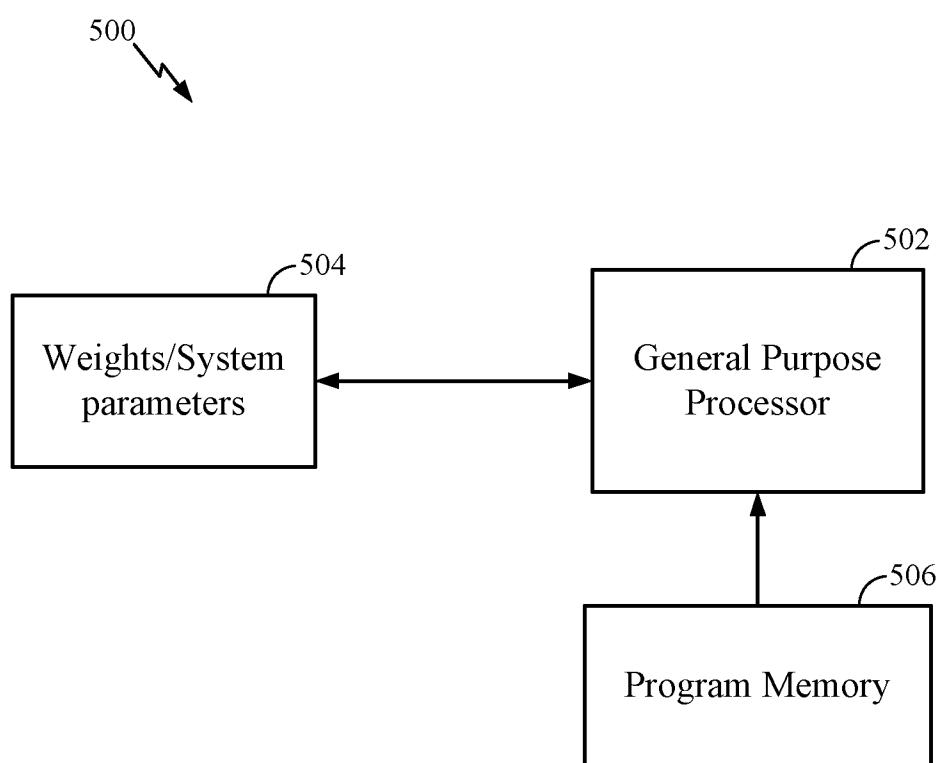
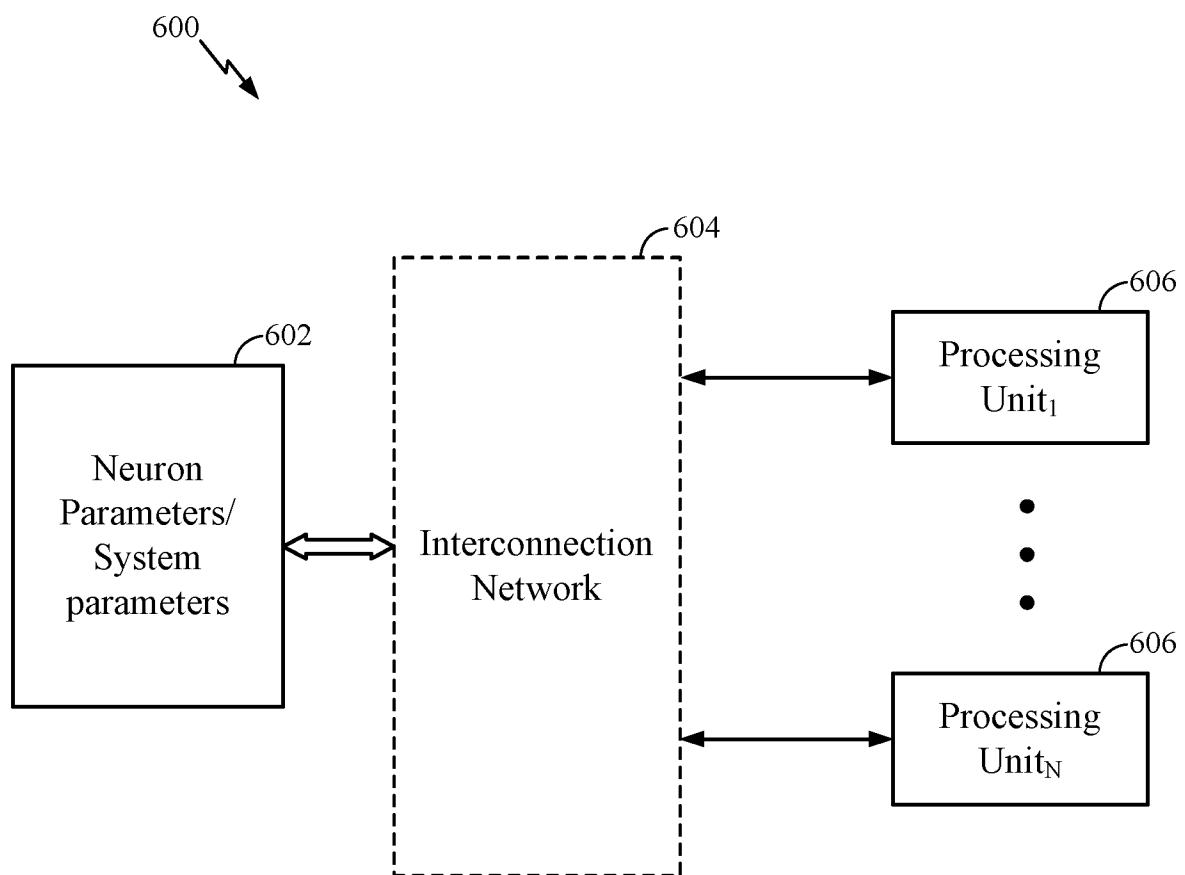


FIG. 5

**FIG. 6**

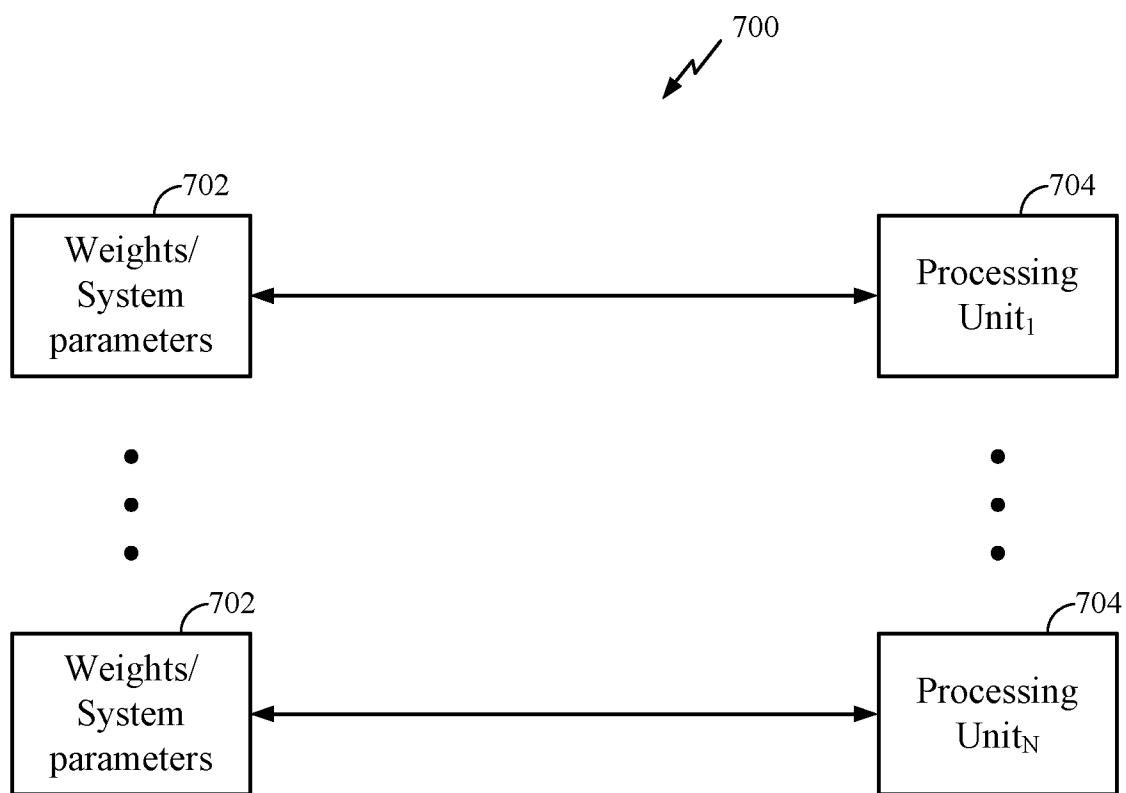
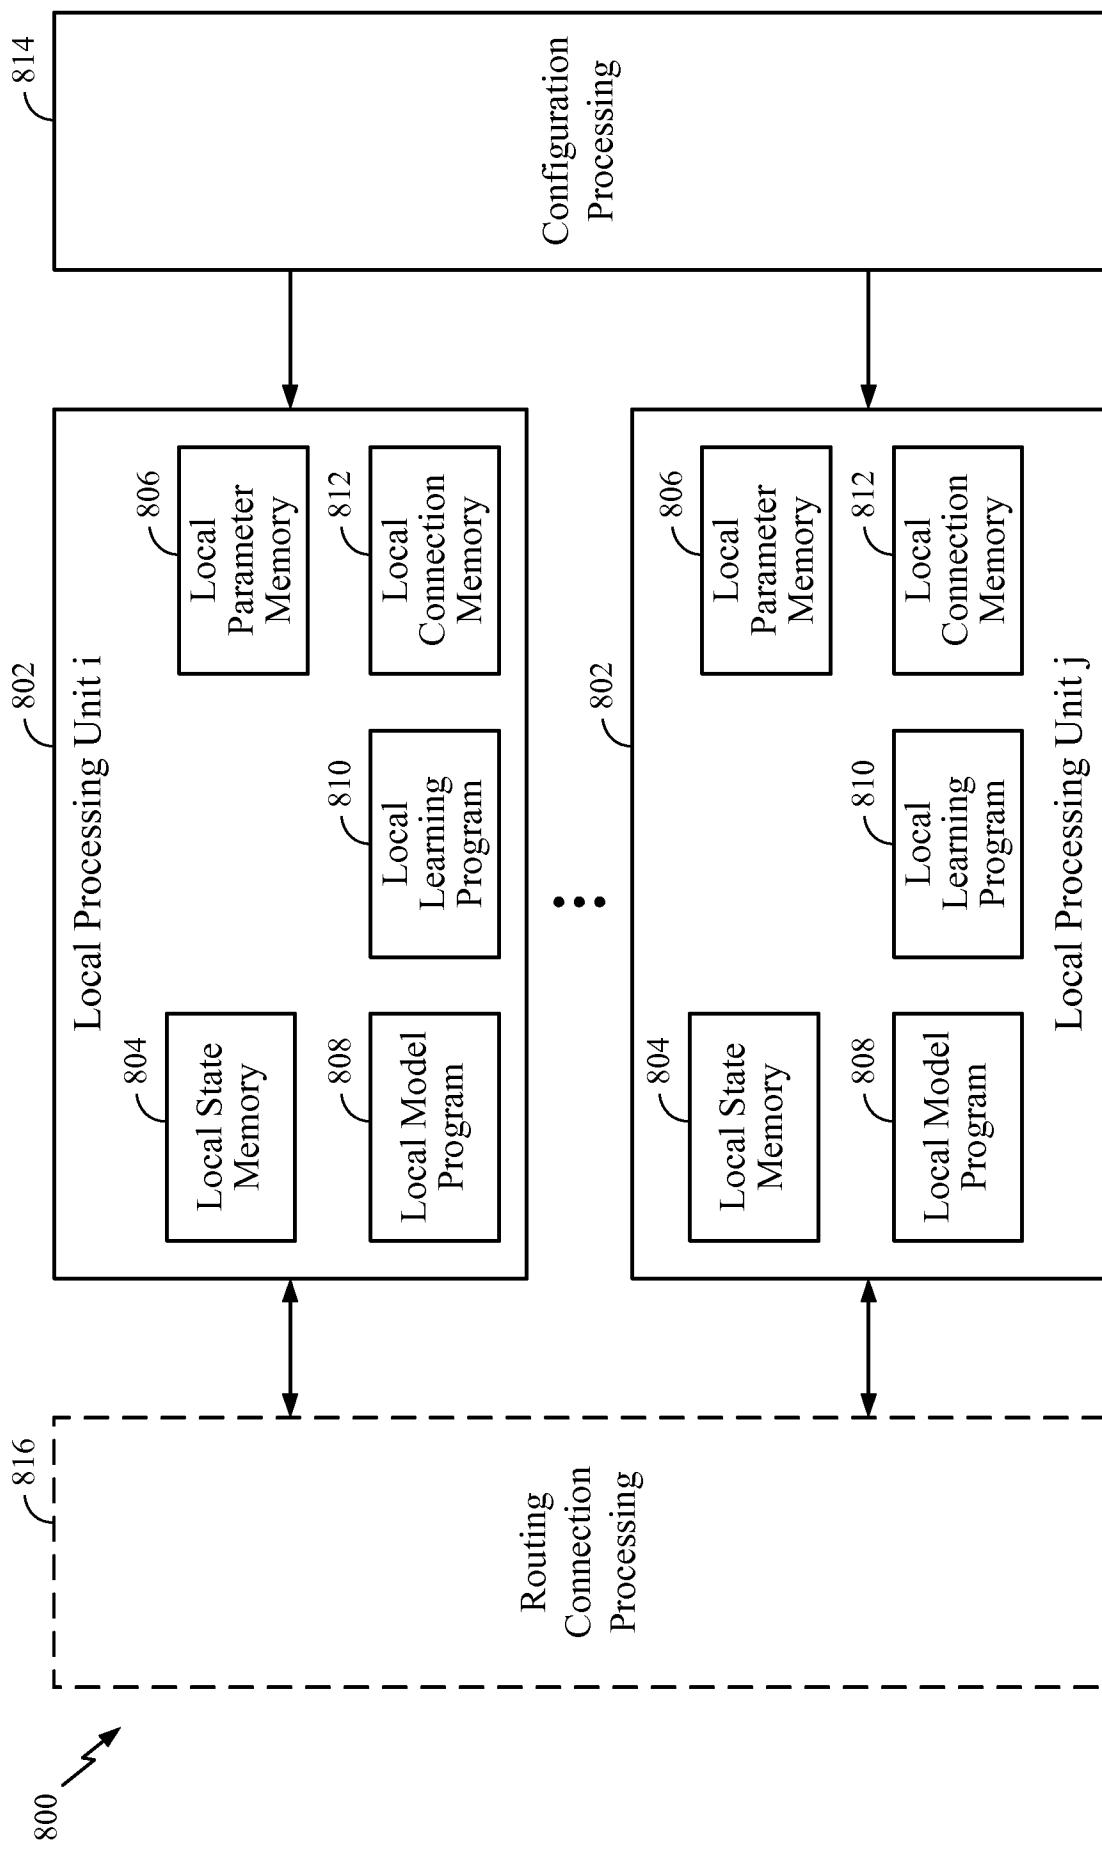


FIG. 7

**FIG. 8**

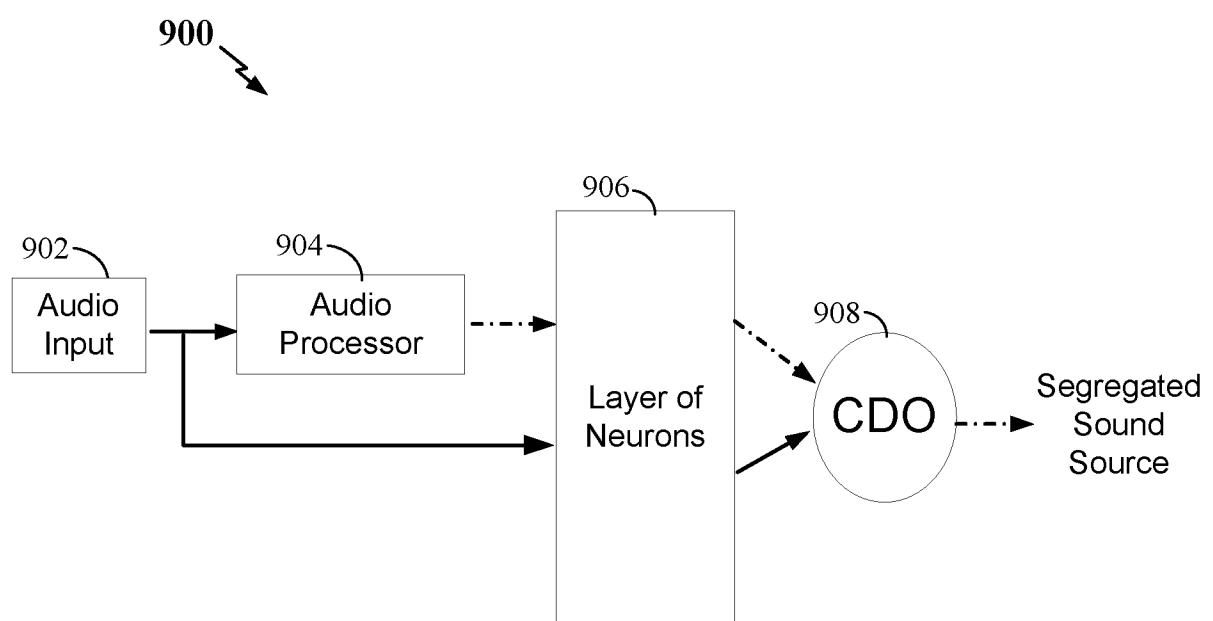
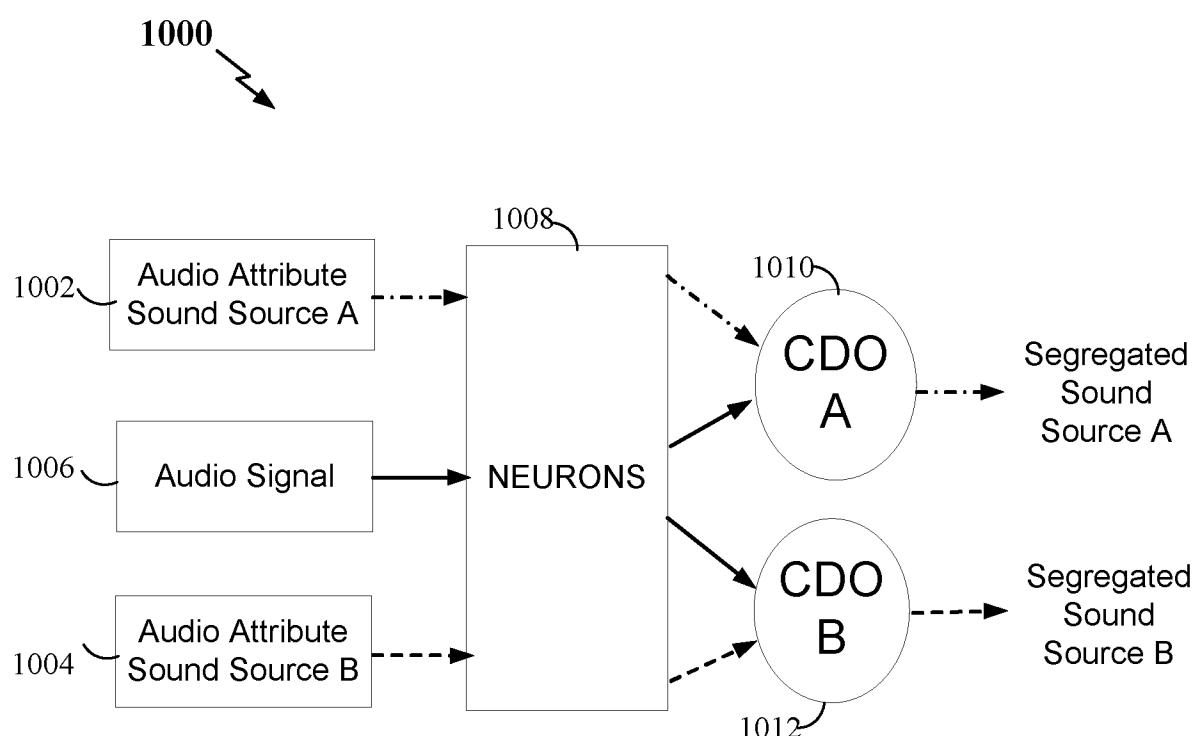
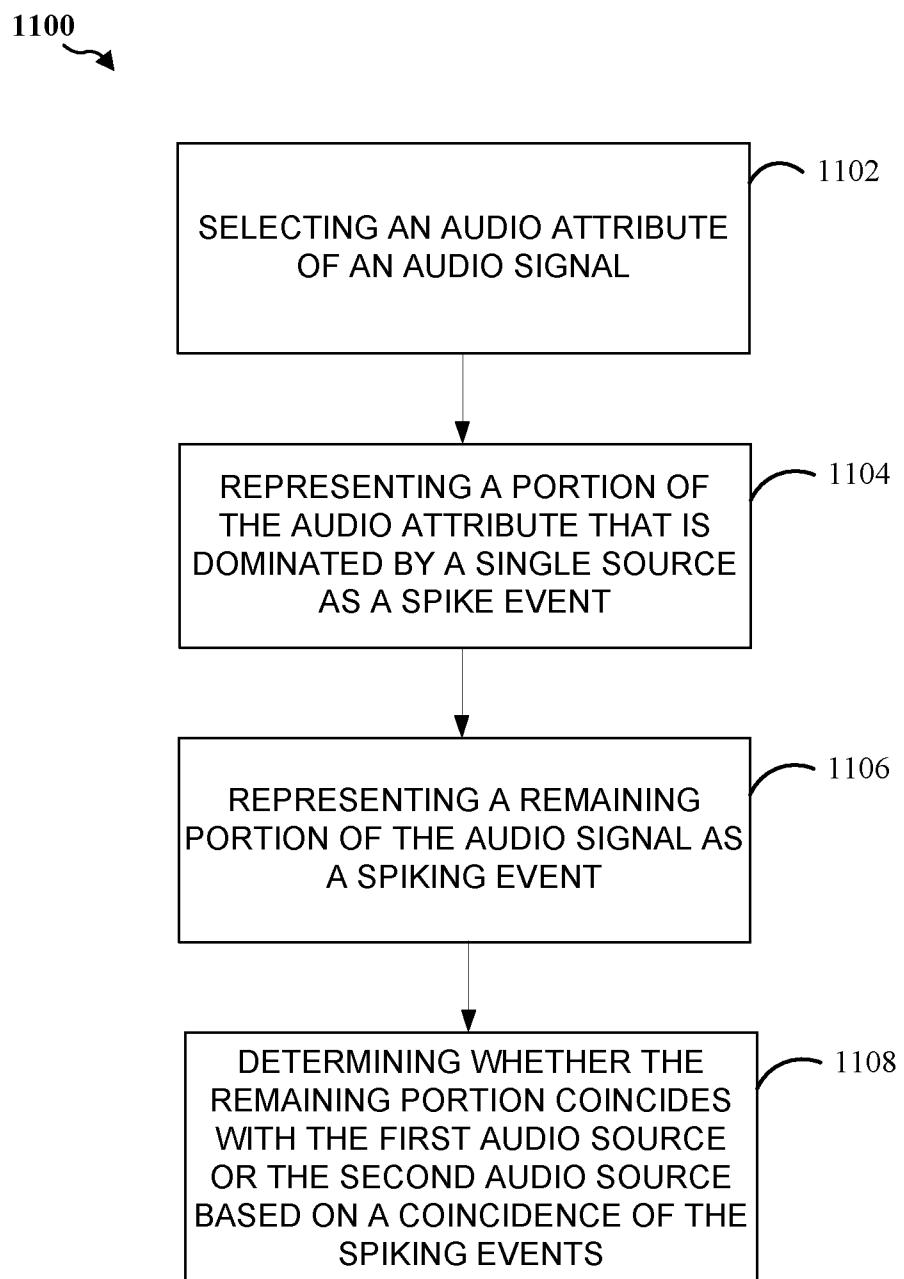


FIG. 9

**FIG. 10**

11/11

**FIG. 11**

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/015697

A. CLASSIFICATION OF SUBJECT MATTER
INV. G06N3/04
ADD. G10L25/30

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)
G06N G10L

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)

EPO-Internal, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	<p>SHIHAB A SHAMMA ET AL: "Behind the scenes of auditory perception", CURRENT OPINION IN NEUROBIOLOGY, vol. 20, no. 3, 22 April 2010 (2010-04-22), pages 361-366, XP028281596, ISSN: 0959-4388, DOI: 10.1016/J.CONB.2010.03.009 [retrieved on 2010-04-02] page 361 - page 365, left-hand column, paragraph 1</p> <p style="text-align: center;">-----</p> <p style="text-align: center;">-/-</p>	1-20

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier application or patent but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search	Date of mailing of the international search report
13 April 2015	23/04/2015
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer Volkmer, Markus

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2015/015697

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	SHIHAB A SHAMMA ET AL: "Temporal coherence and attention in auditory scene analysis", TRENDS IN NEUROSCIENCE, vol. 34, no. 3, March 2011 (2011-03), pages 114-123, XP028172245, ISSN: 0166-2236, DOI: 10.1016/J.TINS.2010.11.002 [retrieved on 2010-11-12] page 114 - page 121, last paragraph -----	1-20
X	MOUNYA ELHILALI ET AL: "Temporal Coherence in the Perceptual Organization and Cortical Representation of Auditory Scenes", NEURON, vol. 61, no. 2, 1 January 2009 (2009-01-01), pages 317-329, XP055182470, ISSN: 0896-6273, DOI: 10.1016/j.neuron.2008.12.005 page 317 - page 328, left-hand column, paragraph 2 -----	1-20
A	ISRAEL NELKEN ET AL: "Encoding Stimulus Information by Spike Numbers and Mean Response Time in Primary Auditory Cortex", JOURNAL OF COMPUTATIONAL NEUROSCIENCE, KLUWER ACADEMIC PUBLISHERS, B0, vol. 19, no. 2, 1 October 2005 (2005-10-01), pages 199-221, XP019208076, ISSN: 1573-6873, DOI: 10.1007/S10827-005-1739-3 page 199 - page 216, right-hand column, paragraph 1 -----	1-20
A	US 2013/204819 A1 (HUNZINGER JASON FRANK [US] ET AL) 8 August 2013 (2013-08-08) the whole document -----	1-20
X, P	LAKSHMI KRISHNAN ET AL: "Segregating Complex Sound Sources through Temporal Coherence", PLOS COMPUTATIONAL BIOLOGY, vol. 10, no. 12, 18 December 2014 (2014-12-18), page e1003985, XP055182467, DOI: 10.1371/journal.pcbi.1003985 page 1 - page 9, right-hand column, paragraph 1 -----	1-20

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No
PCT/US2015/015697

Patent document cited in search report	Publication date	Patent family member(s)		Publication date
US 2013204819	A1 08-08-2013	CN 104094295 A	EP 2812853 A1	08-10-2014