



- (51) International Patent Classification:
G06N 3/02 (2006.01)
- (21) International Application Number:
PCT/US2014/071639
- (22) International Filing Date:
19 December 2014 (19.12.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
14/139,732 23 December 2013 (23.12.2013) US
- (71) Applicant: **QUALCOMM INCORPORATED** (heir of the deceased inventor) [US/US]; ATTN: International IP Administration, 5775 Morehouse Drive, San Diego, California 92121 (US).
- (72) Inventor: **CANOY, Michael-David Nakayoshi**; 5775 Morehouse Drive, San Diego, California 92121-1714 (US).
- (72) Inventor: **HUNZINGER, Jason Frank** (deceased).
- (74) Agents: **LENKIN, Alan M.** et al.; Seyfath Shaw LLP, Suite 3500, 2029 Century Park East, Los Angeles, California 90067-3021 (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, 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.

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

Declarations under Rule 4.17:

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

[Continued on next page]

(54) Title: NEURAL WATCHDOG




FIG. 5

(57) Abstract: A method of monitoring a neural network includes monitoring activity of the neural network. The method also includes detecting a condition based on the activity. The method further includes performing an exception event based on the detected condition.



-
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*
- Published:**
- *without international search report and to be republished upon receipt of that report (Rule 48.2(g))*

NEURAL WATCHDOG

BACKGROUND

Field

[0001] Certain aspects of the present disclosure generally relate to neural system engineering and, more particularly, to systems and methods detecting and recovering from errors in a neural network.

Background

[0002] 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. In some cases, 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. Still, neural networks may fail or experience unexpected behavior. Thus, it is desirable to provide a neuromorphic receiver to detect and recover from errors in a neural network.

SUMMARY

[0003] In one aspect of the presented disclosure, a method of monitoring a neural network is disclosed. The method includes monitoring activity of the neural network. The method also includes detecting a condition based on the activity. The method further includes performing an exception event based on the detected condition.

[0004] Another aspect of the present disclosure is directed to an apparatus including means for monitoring activity of the neural network. The apparatus also includes means for detecting a condition based on the activity. The apparatus further includes means for performing an exception event based on the detected condition.

[0005] In another aspect of the present disclosure, a computer program product for monitoring a neural network having a non-transitory computer-readable medium is disclosed. The computer readable medium has non-transitory program code recorded thereon which, when executed by the processor(s), causes the processor(s) to perform operations of monitoring activity of the neural network. The program code also causes the processor(s) to detect a condition based on the activity. The program code further causes the processor(s) to perform an exception event based on the detected condition.

[0006] Another aspect of the present disclosure is directed to an apparatus for monitoring a neural network having a memory and at least one processor coupled to the memory. The processor(s) is configured to monitor activity of the neural network. The processor(s) is also configured to detect a condition based on the activity. The processor(s) is further configured to perform an exception event based on the detected condition.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] 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.

[0008] FIGURE 1 illustrates an example network of neurons in accordance with certain aspects of the present disclosure.

[0009] 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.

[0010] FIGURE 3 illustrates an example of spike-timing dependent plasticity (STDP) curve in accordance with certain aspects of the present disclosure.

[0011] 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.

[0012] FIGURE 5 is a block diagram illustrating a neural monitor according to an aspect of the present disclosure.

[0013] FIGURE 6 illustrates an example implementation of designing a neural network using a general-purpose processor in accordance with certain aspects of the present disclosure.

[0014] FIGURE 7 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.

[0015] FIGURE 8 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.

[0016] FIGURE 9 illustrates an example implementation of a neural network in accordance with certain aspects of the present disclosure.

[0017] FIGURE 10 is a block diagram illustrating a method for monitoring neural activity according to an aspect of the present disclosure.

[0018] FIGURE 11 is a conceptual data flow diagram illustrating the data flow between different modules/means/components in an exemplary apparatus.

DETAILED DESCRIPTION

[0019] 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.

[0020] 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.

[0021] 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.

[0022] 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

[0023] 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.

[0024] 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). Such behavior can be emulated or simulated in hardware and/or software, including analog and digital implementations such as those described below.

[0025] 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.

[0026] 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 pre-synaptic neurons and neurons of level 106 may be considered post-synaptic 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).

[0027] Biological synapses may be classified as either electrical or chemical. While electrical synapses are used primarily to send excitatory signals, chemical 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.

[0028] 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.

[0029] 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.

[0030] 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.

[0031] 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 204_1 - 204_N (X_1 - X_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 or a voltage, real-valued or 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 206_1 - 206_N (W_1 - W_N), where N may be a total number of input connections of the neuron 202.

[0032] 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, or a voltage, real-valued or 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.

[0033] 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.

[0034] 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

[0035] In hardware and software models of neural networks, the processing of synapse related functions can be based on synaptic type. Synapse types may include 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 execute plasticity functions (or wait 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. Use of varying synapse types may add flexibility and configurability to an artificial neural network.

[0036] There are implications of spike-timing dependent structural plasticity being 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) because 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

[0037] 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.

[0038] 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 post-synaptic neuron's excitation are made even more likely to contribute in the future, whereas inputs that are not the cause of the post-synaptic 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.

[0039] Because a neuron generally produces an output spike when many of its inputs occur within a brief period, (i.e., inputs being sufficiently cumulative 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.

[0040] The STDP learning rule may effectively adapt a synaptic weight of a synapse connecting a pre-synaptic neuron to a post-synaptic neuron as a function of time difference between spike time t_{pre} of the pre-synaptic neuron and spike time t_{post} of the post-synaptic 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 pre-synaptic neuron fires before the post-synaptic neuron), and decrease the synaptic weight (i.e., depress the synapse) if the time difference is negative (the post-synaptic neuron fires before the pre-synaptic neuron).

[0041] 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.

[0042] FIGURE 3 illustrates an exemplary diagram 300 of a synaptic weight change as a function of relative timing of pre-synaptic and post-synaptic spikes in accordance with the STDP. If a pre-synaptic neuron fires before a post-synaptic 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 pre-synaptic and post-synaptic 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.

[0043] 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 of 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 post-synaptic 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

[0044] 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. For example, the arrival time of inputs affects output time and coincidence detection can have a narrow time window. Additionally, 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.

[0045] 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.

[0046] 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 pre-synaptic neuron m to a post-synaptic 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.

[0047] It should be noted that there is a delay from the time when sufficient input to a post-synaptic neuron is established until the time when the post-synaptic 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

[0048] 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.

[0049] 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 (which is different from 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.

[0050] 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.

[0051] 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.

[0052] 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.

[0053] 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 .

[0054] 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.

[0055] 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.

[0056] 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_- .

[0057] 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 specified to reach a particular state. The close form state solutions are:

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

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

[0058] Therefore, the model state may be updated only upon events, such as an input (pre-synaptic spike) or output (post-synaptic spike). Operations may also be performed at any particular time (whether or not there is input or output).

[0059] Moreover, by the momentary coupling principle, the time of a post-synaptic 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:

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

[0060] 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.

[0061] 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.

[0062] 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., 1 ms). This does not necessarily require 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.

NEURAL NETWORK MONITORING SYSTEM

[0063] In some cases, exception conditions may occur in a neural network. The exception conditions may be unexpected instabilities or imbalances that are caused by, for example, network activity, learning functions, plasticity functions, and/or environmental influences. Typical systems may provide controlled debugging environments to debug system failures resulting from an exception condition. Although the debugging environments may be useful to determine the cause of the exception condition, the debugging environments do not mitigate (e.g., prevent) a system failure that may result from an exception condition. Thus, it is desirable to provide a system to detect and/or recover from such conditions in a non-controlled environment, such as the deployed neural network. It should be noted that in the present application an exception condition refers to an event that may potentially lead to a system failure. The exception conditions may be referred to as a condition.

[0064] A typical neural monitor reacts to specified activity. For example, a neural monitor may cut the power to a CPU when the system temperature is above a specified value. One aspect of the present disclosure is directed to a neural monitor that detects instability or imbalance conditions in one or more neural elements and/or processes based on neural activity. That is, the neural monitor may monitor activity patterns of a neural network to determine whether the activity is excepted (e.g., within a range of baseline activity). More specifically, the activity patterns may indicate the potential for a system failure and/or component failure. The neural monitor may be referred to as a probabilistic neural monitor.

[0065] In one configuration, in response to detecting an exception condition, the neural monitor resets and/or reconfigures the primary neural network. Alternatively, the neural monitor may reset and/or reconfigure one or more subsets of the primary neural network. Additionally, the neural monitor may reset the primary neural network and/or one or more subsets of the primary neural network. The neural monitor may reset or reconfigure the software and/or hardware of the primary neural network and/or one or more subsets of the primary neural network.

[0066] In one configuration, the neural monitor is a neural network that receives supervisory input from one or more elements of a primary network. Furthermore, the neural monitor may include individual neural monitors that analyze a single neuron of

the neural network. The individual neural monitors may be collectively used to monitor the activity patterns of the neural network.

[0067] FIGURE 5 illustrates a diagram 500 of a neural processing system 502 that includes a neural monitor 512 according to an aspect of the present disclosure. As shown in FIGURE 5, the neural monitor 512 includes a neural model execution engine 504, an activity monitor 506, and an activity regulator 508. Additionally, a reset block 510 may be coupled to the neural processing system 502. In one configuration, the reset block 510 may be a component of the neural model execution engine 504 or the neural monitor 512. In another configuration, one or more of the activity monitor 506, activity regulator 508, and/or reset block 510 are components of another neural network.

[0068] According to an aspect of the present disclosure, the neural model execution engine 504 is a processing block that executes the neural network (e.g., neural model). The activity monitor 506 monitors the execution of the neural model to detect conditions that specify a response. That is, the activity monitor 506 may monitor the activity of the neural model execution engine 504 and may trigger the activity regulator 508 to regulate (e.g., control) the neural model execution engine 504 to mitigate the detected condition. Alternatively, or in addition, the neural model execution engine 504 may trigger the reset block 510 to reset the neural model execution engine 504.

[0069] As discussed above, the activity monitor 506 may monitor activity of the neural model execution engine 504. In one configuration, the activity monitor 506 determines whether the activity follows a substantially predictable pattern. The activity may include a spiking rate, spiking behaviors/patterns, and/or processing resources. The spiking rate may be monitored for each neuron, for each type of neuron, and/or for each region. Furthermore, the spiking behaviors/patterns may be instantaneous (e.g., absolute) or a rate of spiking. That is, the spiking behaviors/patterns may refer to the distribution of weights in the neural network. Moreover, processing resources activity may include the neural memory bandwidth, processing memory bandwidth, and/or an amount of buffered operations.

[0070] In the present configuration, the activity regulator 508 controls (e.g., corrects) the execution of the neural model execution engine 504 to mitigate the exception condition. In one configuration, the activity regulator 508 reconfigures the neural

model execution engine 504 to mitigate the exception condition. That is, the activity regulator 508 may change the connectivity of the neurons in the neural model execution engine 504. Alternatively, or in addition, the activity regulator 508 may change the number of active neurons and/or change a probabilistic neuron. In another configuration, the activity regulator 508 disables blocks of neurons of the neural model execution engine 504 to mitigate the exception condition. In yet another configuration, the activity regulator 508 applies a correction factor to the neural model execution engine 504 to mitigate the exception condition. In still another aspect, the activity regulator 508 drops neuron spikes. In still a further aspect, the activity regulator 508 reads out network parameters, values, connectivity, and/or history of last time prior to the exception event. In still yet another aspect of the present disclosure, the activity regulator 508 may adjust the gains of neural modulators (e.g., global data) that affect the execution of a neural model. Alternatively, or in addition to, the activity regulator 508 may also adjust the gains to affect the rate of change of system parameters, such as spike generation or internal neuron state

[0071] Additionally, the activity monitor 506 and the activity regulator 508 may control the neural model execution engine 504 to delay a reset. That is, the reset of the neural model execution engine 504 may be undesirable. Therefore, the activity monitor 506 and the activity regulator 508 may work together to reconfigure the neural model execution engine 504 to delay a reset of the neural model execution engine 504.

[0072] Furthermore, in response to a trigger from the activity monitor 506, the reset block 510 may reset some or all of the components of the neural model execution engine 504. The reset block 510 may also trigger a save state to record and/or store the recorded data of the neural model execution engine 504 to a memory unit prior to a reset. The data may include network parameters, values, connectivity, and/or data of a time frame prior to the exception condition. The memory unit may be a component of the neural model execution engine 504 or an external memory.

[0073] FIGURE 6 illustrates an example implementation 600 of the aforementioned neural monitor using a general-purpose processor 602 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 604, while instructions

executed at the general-purpose processor 602 may be loaded from a program memory 606. In an aspect of the present disclosure, the instructions loaded into the general-purpose processor 602 may comprise code for monitoring activity of a neural network, detecting a condition based on the activity, and/or performing an exception event based on the detected condition.

[0074] FIGURE 7 illustrates an example implementation 700 of the aforementioned neural monitor where a memory 702 can be interfaced via an interconnection network 704 with individual (distributed) processing units (neural processors) 706 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, and/or frequency bin information may be stored in the memory 702, and may be loaded from the memory 702 via connection(s) of the interconnection network 704 into each processing unit (neural processor) 706. In an aspect of the present disclosure, the processing unit 706 may be configured to monitor activity of a neural network, detect a condition based on the activity, and/or perform an exception event based on the detected condition.

[0075] FIGURE 8 illustrates an example implementation 800 of the aforementioned neural monitor. As illustrated in FIGURE 8, one memory bank 802 may be directly interfaced with one processing unit 804 of a computational network (neural network). Each memory bank 802 may store variables (neural signals), synaptic weights, and/or system parameters associated with a corresponding processing unit (neural processor) 804 delays, and frequency bin information. In an aspect of the present disclosure, the processing unit 804 may be configured to monitor activity of a neural network, detect a condition based on the activity, and/or perform an exception event based on the detected condition.

[0076] FIGURE 9 illustrates an example implementation of a neural network 900 in accordance with certain aspects of the present disclosure. As illustrated in FIGURE 9, the neural network 900 may have multiple local processing units 902 that may perform various operations of methods described above. Each local processing unit 902 may comprise a local state memory 904 and a local parameter memory 906 that store parameters of the neural network. In addition, the local processing unit 902 may have a local (neuron) model program (LMP) memory 908 for storing a local model program, a

local learning program (LLP) memory 910 for storing a local learning program, and a local connection memory 912. Furthermore, as illustrated in FIGURE 9, each local processing unit 902 may be interfaced with a configuration processor unit 914 for providing configurations for local memories of the local processing unit, and with a routing unit 916 that provide routing between the local processing units 902.

[0077] According to certain aspects of the present disclosure, each local processing unit 902 may be configured to monitor activity of a neural network, detect a condition based on the activity, and/or perform an exception event based on the detected condition.

[0078] FIGURE 10 illustrates a method 1000 for monitoring activity of a neural network. In block 1002, the neural monitor monitors activity of the neural network. Furthermore, in block 1004, the neural monitor detects a condition based on the activity. Finally, in block 1006, the neural monitor performs an exception event based on the detected condition.

[0079] FIGURE 11 is a conceptual data flow diagram illustrating the data flow between different modules/means/components in a neural monitor 1100. The neural monitor 1100 includes a monitoring module 1102 that monitors activity of the neural network. The monitoring module 1102 may monitor activity of an entire neural network or a subset of the neural network. The activity may be received via a connection 1110 with the neural network. The monitoring module 1102 may comprise individual neural monitors that monitor a single neuron of a neural network. The collective individual neural monitors may monitor the entire neural network and/or the subset of the neural network. The monitoring module 1102 may also detect a condition based on the activity. A separate module may be specified for detecting the condition based on the activity (not shown).

[0080] The neural monitor 1100 also includes an exception module 1104 that performs an exception event based on the detected condition. The exception module 1104 receives an input from the monitoring module 1102 when a condition has been detected. The exception module 1104 transmits the exception event via a signal 1112 to the neural network and/or another module to perform the exception event. The exception event may include reconfiguring the neural network or restarting the neural network.

[0081] The apparatus may include additional modules that perform each of the steps of the process in the aforementioned flow chart of FIGURE 10. As such, each step in the aforementioned flow chart FIGURE 10 may be performed by a module and the apparatus may include one or more of those modules. The modules may be one or more hardware components specifically configured to carry out the stated processes/algorithm, implemented by a processor configured to perform the stated processes/algorithm, stored within a computer-readable medium for implementation by a processor, or some combination thereof.

[0082] In one configuration, a neural network, such as the neural monitor of the aspects of the present disclosure, is configured to monitor activity of a neural network. The neural network may include means for monitoring and means for detecting. In one configuration, the means for monitoring and the means for detecting may be the program memory 606, local state memory 904, memory 702, interconnection network 704, processing units 706, processing unit 804, local processing units 902, activity monitor 506, monitoring module 1102, and/or the routing unit 916 configured to perform the functions recited by the monitoring and/or detecting means.

[0083] The neural network may also include means for performing an exception. In one configuration, the means for performing an exception may be the program memory 606, local state memory 904, memory 702, interconnection network 704, processing units 706, processing unit 804, local processing units 902, activity regulator 508, reset block 510, exception module 1104, and/or the routing unit 916 configured to perform the functions recited by the exception means.

[0084] In another configuration, the aforementioned means may be any module or any apparatus configured to perform the functions recited by the aforementioned means. That is, 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 FIGURE 10, those operations may have corresponding counterpart means-plus-function components with similar numbering.

[0085] 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. Further, “determining” may include receiving (e.g., receiving information), accessing (e.g., accessing data in a memory) and the like. In addition, “determining” may include resolving, selecting, choosing, establishing and the like.

[0086] 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.

[0087] 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.

[0088] 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.

[0089] 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.

[0090] The functions described 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.

[0091] 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.

[0092] 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.

[0093] 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.

[0094] 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.

[0095] 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. In addition, 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.

[0096] 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.

[0097] 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.

[0098] 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 monitoring a neural network, comprising:
monitoring activity of a first neural network;
detecting a condition based on the activity; and
performing an exception event based at least in part on the detected condition.
2. The method of claim 1, in which performing the exception event includes correcting execution of the first neural network in response to detecting the condition.
3. The method of claim 1, in which performing the exception event includes reconfiguring blocks of the first neural network.
4. The method of claim 1, in which performing the exception event includes disabling blocks of the first neural network.
5. The method of claim 1, in which performing the exception event includes applying a correction factor.
6. The method of claim 1, in which performing the exception event includes dropping neuron spikes.
7. The method of claim 1, in which performing the exception event includes resetting at least one portion of the first neural network upon detecting the condition.
8. The method of claim 1, in which performing the exception event includes reading out network parameters, values, connectivity, and/or history of last period of time prior to the exception event.
9. The method of claim 1, in which the monitoring, detecting and/or performing the exception event are performed via a second neural network.

10. An apparatus for monitoring neural activity, the apparatus comprising:
a memory unit; and
at least one processor coupled to the memory unit, the at least one processor being configured:

to monitor activity of a first neural network;
to detect a condition based on the activity; and
to perform an exception event based at least in part on the detected condition.

11. The apparatus of claim 10, in which the at least one processor is further configured to correct execution of the first neural network in response to detecting the condition.

12. The apparatus of claim 10, in which the at least one processor is further configured to reconfigure blocks of the first neural network.

13. The apparatus of claim 10, in which the at least one processor is further configured to disable blocks of the first neural network.

14. The apparatus of claim 10, in which the at least one processor is further configured to apply a correction factor.

15. The apparatus of claim 10, in which the at least one processor is further configured to drop neuron spikes.

16. The apparatus of claim 10, in which the at least one processor is further configured to reset at least one portion of the first neural network upon detecting the condition.

17. The apparatus of claim 10, in which the at least one processor is further configured to read out network parameters, values, connectivity, and/or history of last period of time prior to the exception event.

18. The apparatus of claim 10, in which the at least one processor is further configured to monitor, detect, and/or perform the exception event via a second neural network.

19. An apparatus for monitoring a neural network, comprising:
means for monitoring activity of the neural network;
means for detecting a condition based on the activity; and
means for performing an exception event based at least in part on the detected condition.

20. A computer program product for monitoring a neural network, the computer program product comprising:

a non-transitory computer-readable medium having program code recorded thereon, the program code comprising:

program code to monitor activity of the neural network;
program code to detect a condition based on the activity; and
program code to perform an exception event based at least in part on the detected condition.




FIG. 1

2/11




FIG. 2






FIG. 3


4/11

**FIG. 4**


5/11

**FIG. 5**


6/11

**FIG. 6**

7/11

**FIG. 7**

8/11

**FIG. 8**





FIG. 9

10/11

**FIG. 10**




FIG. 11