

Journal Name

Crossmark

PAPER

RECEIVED
dd Month yyyyREVISED
dd Month yyyy

A Neuromorphic Architecture for Scalable Event-Based Control

Yongkang Huo¹⁴, Fulvio Forni¹, Rodolphe Sepulchre¹²³¹The authors are with the University of Cambridge, Department of Engineering, Trumpington Street, CB2 1PZ.² R. Sepulchre is also with KU Leuven, Department of Electrical Engineering (STADIUS), Kasteelpark Arenberg, 10, B-3001 Leuven, Belgium, rodolphe.sepulchre@kuleuven.be.³The research leading to these results has received funding from the European Research Council under the Advanced ERC Grant Agreement SpikyControl n.101054323.⁴ The work of Y. Huo was supported by the UK Engineering and Physical Sciences Research Council (EPSRC) grant 10671447 for the University of Cambridge Centre for Doctoral Training, the Department of Engineering.**E-mail:** yh415@cam.ac.uk, f.forni@eng.cam.ac.uk, rs771@cam.ac.uk**Abstract**

This paper introduces the “rebound Winner-Take-All (RWTA)” motif as the basic element of a scalable neuromorphic control architecture. From the cellular level to the system level, the resulting architecture combines the reliability of discrete computation and the tunability of continuous regulation: it inherits the discrete computation capabilities of winner-take-all state machines and the continuous tuning capabilities of excitable biophysical circuits. The proposed event-based framework addresses continuous rhythmic generation and discrete decision-making in a unified physical modeling language. We illustrate the versatility, robustness, and modularity of the architecture through the nervous system design of a snake robot.

1 Introduction

The divide between the discrete and the continuous pervades today’s control design and technology: “high-level” system design is separated from “low-level” control design, automation is separated from regulation, decision-making is separated from physical interaction. Discrete methodologies use the discrete language of finite state machines while continuous methodologies use the continuous modelling language of differential equations. The resulting need to interface discrete automata and continuous differential equations has made the theory of hybrid and cyberphysical systems a key component of control design. The advantages of separating the discrete and the continuous design are many. Yet, this separation has become a bottleneck in the design of scalable architectures that can control and regulate across a range of spatial and temporal scales [1, 2].

The divide between automation and regulation does not exist in animal machines. Nervous systems make decisions and regulate muscle activation with one and the same architecture that is event-based rather than continuous or discrete. Events occur in physical circuits obeying continuous-time differential equations, yet they can be counted and they can be modelled as the discrete events of a finite state machine [3–5]. Neuromorphic control [4] aims at mixing discrete computation and continuous regulation by borrowing the physical principles and the architecture of excitable neuronal circuits. These neuromorphic approaches build upon four decades of progress in analogue and mixed-signal VLSI that have turned Carver Mead’s vision of “computing with physics” into wafer-scale neuromorphic substrates capable of emulating large-scale neural networks [6–11].

Earlier work in neuromorphic control by the authors has focused on the control and modulation of single neurons [4] or small circuits of few neurons [5]. In contrast, the aim of the present paper is to introduce a *scalable* neuromorphic control architecture for the design of hierarchical event-based machines.

The core element of our methodology is to combine the “rebound excitability” of isolated biophysical neurons with the “Winner-Take-All” paradigm of network computation. The two concepts are classical but, to the best of our knowledge, they have always been studied separately. Rebound excitability is a *cellular* property that has long been identified as a core mechanism of rhythm generation [12]. It has long been studied in neurophysiology and has inspired the design of artificial central pattern generators in rhythmic machines [13–16]. Winner-Take-All computation is

a *network* property at the core of decision-making mechanisms in computational biology and in computer science [17–21]. It has played a central role in the development of neuromorphic computation. A central message of this paper is to show that continuous rhythm generation and discrete decision making become one and same feature when combining the cellular property of rebound excitability and the network property of winner-take-all computation.

We illustrate this principle with the design of a “snake” nervous system via a multi-scale event-based neural network. Figure 1 illustrates the conceptual hierarchical organization of the machine design. This organization is quite generic and the snake example is only chosen for the sake of a concrete illustration. Traditional control architectures would separate the use of physical models for the “low-level” tasks (e.g. muscle actuation) from the use of computational models for the “high-level” tasks (e.g. the supervisory control deciding to move straight or turn left or right). Instead, our design uses the same neuronal architecture all the way from the single neuron to the entire snake controller. Through this unified architecture, we seek to mix “regulation” and “automation” at every stage of the design, so that the resulting event-based machine combines the best of the two worlds, namely the reliability of automation and the tunability of regulation.

The resulting design is inherently modular, event-based, and multi-scale. The modularity stems from the bottom-up design of a hierarchical neural network made of simple functional motifs. The event-based nature of the network stems from the cellular rebound property of the elements. The multi-scale nature stems from the fact that the hierarchical organization of events: events at higher scales are made of ordered sequences of events at lower scales. Such a hierarchy necessarily results in a monotone spread of temporal and spatial scales from finer levels to coarser levels. Because the “tuning” capabilities are always determined at the cellular level, the tuning capabilities are only limited by the finest resolution. Because the “computing” capabilities stem from logical constraints imposed by the network interactions, the control architecture remains reliable at any scale.

The limited scope of the present paper is to focus on the *controller* design only. We do not address the divide between the hardware and the software or between the digital and the analog. Those implementations choices can be made independently from the controller design. The controller is a neuronal circuit, but every part of this circuit can in principle be simulated in a digital computer or emulated in a physical device made of neuromorphic elements.

The rest of the paper is organized as follows: Section 2 reviews the cellular property of rebound excitability in biophysical and neuromorphic circuits. Section 3 introduces the logic gates motifs and Winner-Take-All network. Section 4 introduces the rebound Winner-Take-All network and explains how different patterns can be generated with rebound Winner-Take-All network motifs. Section 5 describes how external events (pulses/biases) tune and interact with these motifs. Section 6 extends the rebound Winner-Take-All network concept from network of cells to network of smaller networks. Section 7 composes the motifs into a five-link snake controller across muscle, pattern, and supervision layers. Discussions and Conclusions are in Sections 8 and 9. All parameter details for all the plots are given in the appendix.




Figure 1: Design requirements for each design level in each task(layer).

Related work. The literature on rhythm generation and Winner-Take-All computation is vast. It is impossible to be exhaustive, but we survey a number of references that directly pertain to the work of the present paper.

Rhythm generation Rhythm generation is a long standing question of neurophysiology rooted in the pioneering concept of central pattern generator (CPG) [12, 22, 23], whose interaction with sensory feedback [24, 25] and neuromodulation [26, 27] generates robust rhythmic activity throughout nervous systems. Computational studies explain how reciprocal inhibition and post-inhibitory rebound (PIR) produce diverse oscillatory modes in the so-called Half-Center Oscillator (HCO) [16, 28, 29]. The biology of rhythms has inspired a lot of research in robotics and neuromorphic engineering, from the early adaptive CPG chip that powered an under-actuated leg [30] to recent spiking implementations on a salamander robot that can swim and walk [13, 14, 31, 32], regulate a pendulum [5], ankle-rehabilitation devices [33], hexapods [34], compliant manipulators [35] and quadrupeds with online CPG optimization [36]. Throughout this literature, the design of the controller *separates* system level computation tasks from CPG rhythm generation, connecting the two with ad-hoc glue logic [37].

Winner-Take-All computation and neuromorphic state machines Winner-Take-All (WTA) networks implement one of the core computational functions of analog computation: they perform analogue *max* of *softmax* computation and serve as the “logic gates” of continuous-time computation [17, 21, 38, 39].

Early continuous-time Hopfield circuits demonstrated *k*-WTA selection [38], and computational complexity studies confirmed the expressive power of WTA modules [17]. Indiveri and colleagues translated these ideas to analogue VLSI, proposing compact soft-WTA arrays for low-power sensing and decision [7, 8, 40]. In spiking hardware, Oster *et al.* showed rapid discrimination under realistic input statistics [39], while Linares-Barranco *et al.* integrated a WTA core with CPGs on a mixed-signal processor [37]. More recently, Cotteret *et al.* embed finite-state machines in attractor-based spiking networks using vector-symbolic arithmetic (VSA) to configure a Hopfield weight matrix, with states encoded in *rate* space [41, 42].

Beyond hardware, a number of studies have enabled computational functions that are essential in human high-level cognitive systems to be implemented on spiking devices, such as finite state machines, auto associative memories [40–44] and logic gates, to be directly implemented in recurrent spiking dynamics.

Dynamical winner-take-all networks. The concept of generating dynamical spatio-temporal patterns from winner-take-all principles also has a long history, best illustrated by the principle of

winnerless competition in neural dynamics [45]. Winnerless competition circuits interconnect *rate* neuronal models to build heteroclinic trajectories between saddle points, which raises challenges in ensuring both sensitivity to inputs and robustness against noise [45, 46]. A number of studies try to resolve this issue by determining the condition of robust heteroclinic [47–52] or adding external controllers [53]. Previous work has shown that winnerless competition controllers deliver tunable patterned motion, including snake locomotion [54].

The key difference between the mechanism of winnerless competition and the approach of the present paper is that we separate the *cellular* property of rebound excitability and the *network* property of Winner-Take-All computation. Rhythm generation in neuronal networks that have no excitability property at the cellular level requires an accurate tuning of the synaptic interconnections. In contrast, the restricted role of synaptic interactions in rebound-Winner-Take-All networks is to logically constrain the sequence of events. The decoupling between “event generation” and “event orchestration” is a key source of robustness and modularity in the proposed architecture.

2 Cellular Level: Rebound Excitability, Neuronal Models, and Synapses.

2.1 Rebound Excitability

The mechanism of rebound excitability has been studied in great detail since the early days of neurophysiology. We refer the reader to [4] for a detailed survey of the relevant physiology as well as the neuromorphic question of designing artificial neurons with tunable rebound excitability, we refer the reader to [4].

In this paper, we rely on the classical circuit architecture laid out in the pioneering work of Hodgkin and Huxley [55]. A neuron is modeled as an electrical circuit that connects one leaky capacitor (the passive membrane) to a parallel bank of active filters defined by the series combination of a battery and a memristive element, as shown in Figure 2. Each memristive element obeys the Ohmic law $i = gv$, where the ‘conductance’ g depends nonlinearly and dynamically on the membrane voltage.

When the conductance depends solely on the internal capacitor voltage, v , the active filter represents an ion channel (e.g., sodium, potassium, calcium). This dependency defines an *internal conductance*, regulating ion flow through the membrane without external neuronal influence. If the conductance, however, depends on the membrane voltage of an external neuron, it is classified as an *external conductance* and models a synaptic interconnection. Whether internal (ion channel) or external (synapse), the precise model of a conductance g varies across the literature. A conductance is often characterized by its activation and inactivation functions, each characterized by a specific voltage range and time-constant.




Figure 2: Circuit model of neuron

Excitability is the key property of a biophysical neuron [56]: the voltage response to a current stimulus is passive for small stimuli and *event-based* when the stimulus exceeds a threshold. The simplest event-based response of a neuron is a single spike. An important property of biophysical neurons is that both depolarizing inputs (excitation - the input current flow increases the membrane voltage) and hyperpolarizing inputs (inhibition - the input current flow decreases the membrane voltage) can trigger an event. In the latter case, the event occurs at the end of the inhibition stimulus. Such an event, illustrated in Figure 3a, is called a *rebound event*, and the corresponding excitability property is called *rebound excitability*.

In conductance-based models, the simplest spike event is the result of the interaction of two currents: a fast-activation inward current (sodium channels, in the model of Hodgkin and Huxley) that creates positive feedback, and a slow-activation outward current (potassium channels, in the model of Hodgkin and Huxley) that repolarizes the neuron through negative feedback. Additional internal conductances can be used to control the excitability features of the neuron, including its sensitivity and robustness to dynamic uncertainties and perturbations. These conductances also enable the generation of specific event types. In this paper, we limit our discussion to spike and burst events at the cellular level.

Rebound excitability is a crucial mechanism for sequentially propagating events through a chain of neurons connected by inhibitory synapses. An event in a presynaptic neuron creates transient inhibition in the postsynaptic neuron, which subsequently terminates in a rebound event. This basic mechanism enables the generation of reliable sequences of events. The duration and properties of each event are tuned by the internal conductances of the respective neuron; a cellular-level tuning mechanism. The overall sequence of these events is instead defined by the synaptic interconnection; a network-level tuning mechanism.

2.2 Rebound Spiking and Bursting

Throughout the paper, we focus on qualitative properties that do not depend on fine details of conductance models. Therefore, for simplicity, we adopt the model proposed in [4]. This model captures the essential properties of neurons, including *rebound excitability*.

As shown in (1), the membrane voltage (RC circuit) of a *rebound spiking neuron* is driven by the external current I_{ext} and by two internal currents: a fast depolarizing current I_f (static nonlinearity) and a slow repolarizing current I_s (cascade of low-pass filter and static nonlinearity). The parameters α s regulate the strength of the currents. $\bar{V}$ models the presence of batteries. The rebound spiking excitability of the neuron is illustrated in Figure 3a.

$$\begin{aligned} C\dot{V} &= -RV - I_f - I_s + I_{\text{ext}} \\ I_f &= -\alpha_f^- \tanh(V - \bar{V}) \\ I_s &= \alpha_s^+ \tanh(V_s - \bar{V}) \\ \tau_s \dot{V}_s &= -V_s + V \end{aligned} \tag{1}$$

In a similar way, the mathematical model of a *rebound bursting neuron* is given (2), characterized by the introduction of a new pair of currents, at a slower time scale (an additional term within the definition of I_s and the ultra-slow current I_{us}). This is the time-scale of the burst, a cluster of spikes separated by periods of quiescence (no firing). The role of these additional terms is to produce a slower overshoot, similar to a long smoothed spike, which then turns into a burst. The rebound bursting excitability of the neuron is illustrated in Figure 3b, which shows the membrane voltages and the ion currents at 4 stages of rebound excitation.

$$\begin{aligned} C\dot{V} &= -RV - I_f - I_s - I_{us} + I_{\text{ext}} \\ I_f &= -\alpha_f^- \tanh(V - \bar{V}_f) \\ I_s &= \alpha_s^+ \tanh(V_s - \bar{V}_f) - \alpha_s^- \tanh(V_s + \bar{V}_s) \\ I_{us} &= \alpha_{us}^+ \tanh(V_{us} - \bar{V}_s) \\ \tau_s \dot{V}_s &= -V_s + V \\ \tau_{us} \dot{V}_{us} &= -V_{us} + V \end{aligned} \tag{2}$$

As before, the parameters α s regulate the strength of the currents and $\bar{V}_f$ and $\bar{V}_s$ model the effect of batteries. As shown in [4], a burst rebound event is similar to a spike rebound event but it has a well-defined and tunable duration. The burst duration is regulated by the ultra-slow parameter α_{us}^+ , as shown in Figure 4. Hence we can regard a spiking event as an event of *fixed* duration and a bursting event as an event of *variable* duration.




Figure 3: Rebound neurons.




Figure 4: Two voltage trajectories, A and B, of the same neuron for different values of α_{us}^+ . The input current is an inhibitory step from -1.5 to -5 with duration 50 ms. Trajectory A (B) corresponds to a smaller (larger) α_{us}^+ .

2.3 Inhibitory and Excitatory Synapses

Synaptic currents model directional interconnections between neurons. A synaptic current is (dynamically) determined by the presynaptic neuron membrane voltage and affects the postsynaptic neuron membrane voltage. In this paper we adopt the simplified synaptic model of [4]. A synaptic current I_{ij} of type k from the presynaptic neuron j (voltage V_j) to the postsynaptic neuron i has the expression

$$I_{ij} = G_{ij} \sigma(V_{ij} - \bar{V}_{ij})^k$$

$$\sigma(V_{ij} - \bar{V}_{ij}) = \frac{1}{1 + e^{-\alpha_{ij}(V_{ij} - \bar{V}_{ij})}} \quad (3)$$

$$\tau_{ij} \dot{V}_{ij} = -V_{ij} + V_j,$$

where G_{ij} is the synaptic weight, $\bar{V}_{ij}$ represents the contribution of the battery, and α_{ij} regulates the slope of the (sigmoidal) function $\sigma(\cdot - \bar{V}_{ij})$ in the neighborhood of $\bar{V}_{ij}$. As for the case of internal currents, synaptic currents are modeled as the cascade of a low-pass filter with time constant τ_{ij} and a biased sigmoidal nonlinearity. The synaptic weight G_{ij} is *positive* for *excitatory* synapses and *negative* for *inhibitory* synapses.

The synaptic model of (3) simplifies the classical Ohmic model of a physiological synapse. The latter has an additional multiplicative factor that determines the sign of the current. In contrast, in (3), the excitatory or inhibitory nature of a synapse is fully determined by the sign of G_{ij} ; an easier characterization. The synapse can be thought of as a nonlinear filter on the presynaptic voltage. The quality of this filtering action provides different ‘readouts’ of presynaptic events. For example,

a fast synapse with a high $\bar{V}_{ij}$ driven by a bursting neuron will produce a cluster of current pulses (several events in current space). In contrast, a slow synapse with a low $\bar{V}_{ij}$ will low-pass filter the cluster of spikes into a slowly rising and decaying current bump (a single event in current space).

In this paper, all synapses are confined to three time-scales: fast ($\tau \simeq 0.5$), slow ($\tau \simeq 50$), and ultra-slow ($\tau \simeq 100$). In all figures of this paper, fast synapses are represented with blue, slow synapses are represented with red and ultra-slow synapses are represented with black. We use an *arrow* for *excitatory* synapses and a *circle* for *inhibitory* synapses.

3 Functional Level: Decision Making

Logical operations such as NOT, AND and OR are the fundamental elements of decision making. In this section, we provide a neuromorphic realization of these classical logical elements. We then move to a more complex network, illustrating a neuromorphic realization of the classical Winner-Take-All mechanisms (WTA) [21]. Logic gates and winner-take-all mechanism realizations are illustrated in Figure 5.

NOT gate: A single inhibitory synapse implements a NOT gate. The NOT function means that an event in the presynaptic neuron forbids a simultaneous event in the postsynaptic neuron.

AND gate: Two rebound neurons (input nodes) connect with excitatory synapses to a third rebound neuron (output node). The maximum gain of synapses can be chosen such that events in the two presynaptic neurons are necessary to trigger an event in the postsynaptic neuron.

OR gate: The architecture is similar to the AND gate. However, the maximum gain of the synapses is increased to guarantee activation of the postsynaptic neuron whenever either one of the two presynaptic neurons spikes.

In digital electronics, OR (AND) gates can also be realized using NOT and AND (OR) gates. Similar constructs can be proposed in our neuromorphic setting. These derived neuromorphic circuits would exhibit longer propagation delays but offer interesting features in terms of robustness and tunability of their response to a stimulus. For instance, the OR gate proposed in Figure 5 shows sensitivity to the selection of synapses weights. In contrast, though consuming more energy and introducing greater delays, its realization through NOT and AND neuromorphic circuits offers reduced sensitivity to the specific choice of synaptic gains in the network.

Winner-Take-All (WTA) network: This is a competitive neural circuit that compares multiple inputs and amplifies the strongest one while suppressing the rest, typically through lateral inhibition. Functionally, it converts small differences in input strength into a clear selection. Winner-Take-All network is typically an all-to-all inhibitory network, as shown at the bottom of Figure 5. To avoid clutter, the all-to-all connectivity is represented by a central node (see also Figure 6).

The raster plot on the right of the WTA network in Figure 5 illustrates its input-output behavior. The inputs u_1 - u_3 are the currents injected into neurons below 1-3. Neuron 1 receives the highest input at the start while all other neurons are simply affected by noise. Hence, neuron 1 becomes the winner. Even though the input of neuron 2 is later increased, the selected ‘winner’ neuron does not change. This demonstrates the hysteresis behavior at the core of the Winner-Take-All network. The network becomes again sensitive to external inputs only when the input to neuron 1 is reduced. As shown in Figure 5, in such a case a slightly higher input to neuron 3 than to neuron 2 sets the new winner (neuron 3).




Figure 5: Decision Making Motifs. Fast synapses ($\tau = 0.5$) are represented with blue, slow synapses ($\tau = 50$) are represented with red. Arrow for excitatory connection and circle for inhibitory connections. Dashed arrows denote input and output connections.




Figure 6: Equivalent representations of an all-to-all inhibitory network.

4 Functional Level: Pattern Generation

Unlike pure logic operations, events in the real world unfold in time, and neural circuits respond to them with specific time dynamics. This dynamic time response is crucial to shape how the circuit interacts with its external environment. In this section, we explore this temporal dimension by introducing rebound neural networks for pattern generation. We proceed from the simplest case, the Half-Center Oscillator, to the more general, the Rebound Winner-Take-All network, showing how rebound at the cellular level and inhibitory competition at the network level combine to yield robust spatio-temporal patterns.

Half-Center Oscillator (HCO): It consists of two rebound neurons reciprocally connected by inhibitory synapses, as shown at the bottom row of Figure 7. It is the simplest pattern generator built from rebound neurons [12].

A useful way to understand the network is to start with a *single* inhibitory synapse from neuron 1 to neuron 2 (Figure 7, top row). When neuron 1 spikes, it inhibits neuron 2; upon release of inhibition, neuron 2 produces a rebound spike. With *reciprocal* inhibition, this rebound triggers the reverse action, yielding an alternating left-right pattern. This rebound-based mechanism is classical and extensively documented [26, 27].

In our implementation, the synapse dynamics for the HCO use $\tau = 0.5$, $\bar{V} = -1$, $\alpha = 50$, and weight -2 . The threshold is chosen near or below the neuron's spiking threshold so that the postsynaptic neuron activates only when the presynaptic neuron has returned to rest; τ is fast enough to preserve the spike after filtering; α is large enough to saturate during rebound. The resulting behavior is shown in Figure 7.




Figure 7: The half-center oscillator is the simplest example of a RWTA network. The two neurons of the network rebound in alternance.

Rebound Winner-Take-All (RWTA): It consists of network of synaptically interconnected rebound neurons characterized by the presence of fast all-to-all inhibitory connections to enforce that at most one neuron can spike at any time (the *single-winner* constraint). The Rebound Winner-Take-All network generalizes the functionality of the HCO to networks of rebound neurons of arbitrary size. Figure 14 illustrates 3 examples of RWTA motifs. They are: a ring oscillator of spiking neurons, a ring oscillator with spiking and bursting neurons, and a dual ring oscillator.

The topology of these networks is based on two complementary principles: (i) *Fast inhibition* for ‘state restriction’. The all-to-all inhibition limits the event state space to n states for n neurons. At any time, exactly one neuron can be ‘on’ while the rest must be ‘off’. (ii) *Slow excitation* to ‘enable specific transitions’. Directed slow excitatory couplings bias the competition so that, at the end of an event, a designated postsynaptic neuron is favored as the next winner. The time-scale separation between inhibition and excitation is a key feature. Excitatory synapses are *slow* relative to inhibitory synapses so that the bias persists through the event and into the rebound stage of the next neuron. The following networks show how to construct specific temporal patterns and also illustrate the flexibility of the Rebound Winner-Take-All network.

Spiking Ring Oscillator: A ring of N spiking neurons with fast all-to-all inhibition and slow excitation connections forming a ring produces an autonomous traveling wave $1 \rightarrow 2 \rightarrow \dots \rightarrow N$, in

the absence of external input. The length of the inter spike interval is given by the combined contribution of spike interval and rebound interval (see Figure 3a). The former is mostly affected by the time constant of the fast inhibitory synapses. The latter is sensitive to the time constant and weight of the slow excitatory synapses and it also increases monotonically with the input magnitude. More details can be found in [57].

Notably, for the case of two neurons $N = 2$, the RWTA network corresponds to the Half-Center oscillator. This also shows that the addition of excitatory connection on the HCO network does not violate its Winner-Take-All characteristics.

Spiking and Bursting Ring Oscillator: The features of the RWTA network are robust to heterogeneity. Replacing a spiking neuron with a bursting neuron does not affect the pattern generation abilities of the network. The effect of the bursting neuron is to stretch inter-wave waiting interval while preserving the same sequence of events. Only the differences in intrinsic time-scales and thresholds might need to be compensated by neuron-specific bias currents and appropriately matched synaptic time constants. In our simulation, we do not introduce any compensation, that is, the interconnection parameters are kept constant between the spiking and the bursting cases. This shows one of the advantages of decoupling the network design from the cellular design. The topology of the network only controls the sequencing of events, but the events themselves are neuronal properties. This decoupled role of the network topology and the neuronal parameters is what makes the network topology robust to cellular heterogeneity.

Multiple Ring Oscillators as Pattern Selector: This example demonstrates how the synapses determine the rhythmic properties of the network. Specifically, connecting two rings that *share a common inhibitory hub* turns the network into a state machine. This architecture enables the network to select between two different patterns based on the external input. Because the hub enforces single-winner behavior across rings, a brief external pulse to any neuron can switch the active pattern. Figures 14 illustrates toggling between two rhythmic patterns, demonstrating discrete state transitions on top of rhythmic generation.

Across all these examples, we can identify the following general properties of RWTA networks:

- P1 Single-winner:** at any time, at most one neuron is in an event state (the winner).
- P2 Refractoriness:** events end with a refractory period, preventing a neuron from winning twice in rapid succession.
- P3 Competitive selection:** during rebound, all neurons compete; slow excitation and/or depolarizing inputs designate the next winner.
- P4 Event-driven sequencing:** cycles of slow excitation yield endogenous cyclic event sequences (e.g., ring waves).

The specific separation of inhibition and excitation is at the core of these properties. In the RWTA architecture, the design of the states (governed by fast inhibition) is effectively decoupled from the design of the transitions (governed by slow excitation). This qualitative interconnection design, which relies on topology and time-scales rather than on precise values, ensures robustness against parameter variability and neuron mismatch.

We complete this section by emphasizing that the computational principles of our neuromorphic architecture are consistent with the biophysical architecture of EI networks studied in neuroscience [58, 59]. An EI network of rebound neurons was previously studied in [60] to show how cellular neuromodulation can easily reconfigure the network *states*. A neuromorphic neuron in our architecture corresponds to a pair of excitatory and inhibitory neurons in EI networks. The inhibitory layer must be densely connected to approximate the all-to-all topology of a rebound network, and the inhibitory synapses are typically of the GABA type. In contrast, the excitatory layer is sparsely connected by AMPA synapses. While AMPA synapses are normally faster than GABA synapses, our *slow* excitatory synapse can be regarded as a model of the synaptic delay resulting from event propagation through several nodes in the excitatory layer.




Figure 8: Rhythmic states generated by slow cyclic excitatory connections in a RWTA network.

5 Functional Level: Interaction of the RWTA Network

The rebound-Winner-Take-All architecture encodes reliable endogenous *states* (i.e. event sequences that do not require external input), but its *rebound* nature makes it also eminently interactive with the external environment. In the absence of external inputs, the rhythm is endogenous, that is, only determined by the internal network topology. But external inputs can either re-orchestrate or tune the rhythm as follows:

Phase switching (*local external interaction*): A *brief pulse* $u_i(t)$ applied to neuron i during its rebound stage raises its membrane voltage V_i , relative to its competitors. If the pulse arrives before any neuron has crossed threshold, i wins the current competition and the network *jumps* to phase i . If the pulse arrives during inhibition, the neuron will have a higher initial voltage at the start of next rebound and hence will increase the probability to win the next round. Thus, the RWTA network functions as an event-driven Winner-Take-All selector that can be either hard-reset or synchronized to external events. This is illustrated in Figure 9. Neuron 2 receives pulses at two different times and each time the network pattern is reset to Neuron 2.

Frequency modulation (*global external control*): A *slow modulation of the common bias* represented by $u_i \mapsto u_i + \Delta u$ for all neurons i , increases the rebound speed of *every* neuron. This stretches or compresses the rebound interval of each neuron while leaving the firing order unaffected. As a result, the RWTA network behaves as a current-controlled rhythmic pattern generator whose frequency can be tuned continuously either in open loop (constant Δu) or in closed loop (Δu is continuously regulated by a feedback loop). Figure 9 shows how a hyperpolarization step change in the the average current input decreases the frequency of the generated rhythm.

These two cases show how the pattern generation capabilities of the RWTA network exhibit features from both discrete localized event-based interaction and continuous global neuromodulation behaviors, resulting in a versatile yet robust pattern generator.




Figure 9: Two distinct mechanisms of external control: fast and localized external events can phase advance or reset the endogenous pattern; slow and global external inputs modulate global properties of the rhythm, e.g. its frequency.

6 Multiscale States and Network Hierarchy

In this section we show that the rebound Winner-take-all network is hierarchical and multi-scale: neurons generate events, networks generate patterns of events that can themselves be regarded as new events at a broader scale, and so on. We limit our discussion to a few specific but conceptual examples.

Figure 10 shows the architecture of a RWTA of RWTA networks. The top row represents the structure and simulation of a three nodes RWTA network. Rhythmic sequences are turned into events (of events) by adding ultra-slow all-to-all inhibitory synapses that control that terminate the cyclic pattern after a fixed duration. The effect of the ultra-slow inhibitory connections is illustrated in the top row simulations of Figure 10. All the neurons are initially driven by a strong external inhibition, released at, which is released at 200ms. Right after the release, the network rebounds, generating a rhythmic pattern. In contrast to the previous section, the pattern terminates after a few iterations. In other words, the pattern itself becomes a new event made of a sequence of events at finer scale.

In the second row of Figure 10 we have constructed a HCO of rebound networks. The two rebound networks are connected by slow all-to-all synapses, represented by a purple node. Their time scale is faster than the slow excitatory synapses (red) but can be slower or equal to the time scale of the fast inhibitory synapses internal to each network (blue). The simulation illustrates a clear rhythmic alternating pattern between the patterns of each rebound networks.

As last example, we create a three-node RWTA network of RWTA networks, as shown in the bottom row of Figure 10. The outer layer of excitatory connections between each node (black) are ultra-slow synapses with the same time constant as the all-to-all ultra-slow inhibitory synapses internal to each node (also in black). The simulation illustrate the pattern of a 3-ring oscillator of 3-ring oscillators.

These examples suggest that the design of a RWTA network is inherently hierarchical. Specific sequences of events form new events. With the same interconnection rules, those new events can be sequenced to create new patterns, etc.. The hierarchy of time-scales and the complementary role of excitation and inhibition are key to enable the hierarchy of the architecture. These example explore a two-layer network but the quality of the result suggest scalability towards more complicated hierarchies.




Figure 10: The RWTA architecture is inherently hierarchical, owing to the interlacing of E and I connections and the monotone ordering of time-scales.

7 Neuromorphic control of a five-link robotic snake

7.1 Snake mechanics, muscle actuation, and neuromorphic controller basics

We illustrate the flexibility and reliability of the proposed neuromorphic architecture by constructing a motion controller for a simplified, five-link snake robot. Our goal is to show that the functional network motifs presented in the previous section can be easily composed to build complex behavior across different scales. The illustrative purpose of this section is to suggest that the proposed architecture enables the design of a simple ‘nervous system’ capable of driving the robot’s locomotion.

The plant: the snake ‘skeleton’ consists of five rigid segments connected by four revolute joints. Each joint is actuated by a pair of ‘muscles’. Each muscle is modeled by (i) a torque spring with a variable reference point, (ii) a linear torque spring, and (iii) a linear damper [61, 62], as shown in Figure 11. The contraction/relax of a muscle is modeled via the change in the reference trajectory of the first spring. In what follows, we regard the combination of snake mechanics and muscles as the *plant*.




Figure 11: Plant: The five-link snake and muscle model

The controller: our neuromorphic architecture is used to design a hierarchical *controller*. The controller outputs are the muscle reference trajectories r on each side of each joint. As summarized in Figure 1, the controller is organised into three layers: (i) Muscle Actuation Layer (local joint actuation) (ii) Coordination Layer (gait phasing and traveling waves) (iii) Supervisory Control Layer (task-level switching). All three layers use the same RWTA architecture. They differ only in topology and time-scale. Each layer is further organized into event, functional, and system levels.

The controller architecture uses five different categories of synapses: inhibitory (I), slow inhibitory (I_s), slow excitatory (E), fast high-threshold excitatory (E_f), fast low-threshold excitatory (E_{fl}), and ultra-slow excitatory (E_{us}). Their parameters are listed in Table 1. Fast ($\tau = 0.5$), slow ($\tau = 20$), and ultra-slow ($\tau = 100$) synapses are coloured blue/red/black in the figures. Arrows denote excitation and hollow circles denote inhibition.

Table 1: Parameters of synapses.

Synapse type	τ	α	$\bar{V}$
I/E_{fl}	0.5	50	-2.5
E/I_s	20	50	-2.5
E_f	0.5	50	0
E_{us}	100	50	-2.75

7.2 Muscle Actuation Layer

Goal. Produce alternating left/right muscle activation at each joint, for local swinging.

Event level. At this level, the basic motor command is a *muscle contraction event*. We encode a contraction on one side of a joint as a *burst*: a rebound burst drives one of the two sides ON for a given duration, while the other side is suppressed. Because the burst duration is set by the ultra-slow current of the neuron, the duration is directly tunable.

Functional level. Each joint is driven by a *half-center oscillator* (HCO), given by two rebound bursting neurons connected by fast reciprocal inhibition I . This produces left/right alternation of the muscle activation of each joint. Each neuron in the HCO connects via a slow excitatory synapse E to the corresponding muscle reference point r (Figure 12). The HCO block acts as a local muscle controller: one neuron excites the left muscle, and the other neuron excites the right muscle.

System level. We have four HCOs, one per joint. Together with the muscles, they form **Muscle Actuation Layer Block**, represented in Figure 12. The coordination of these four HCOs, that is, the relative phase of the HCOs, is essential for locomotion. Different ‘initial conditions’ produce distinct trajectories. This is shown in Figures 13a and 13c show two trajectories generated by two different ‘initial conditions’. The corresponding membrane traces are reported in Figures 13b and 13d, respectively. The coordination among the four HCOs can be regulated. This is the role of the next layer.




Figure 12: Muscle Actuation Layer. Each joint is controlled by a Half-Center Oscillator (HCO): two mutually inhibitory rebound neurons that alternate activity and drive the muscles via slow excitatory synapses E . Four HCOs (one per joint) form the Muscle Actuation Layer Block.




Figure 13: Muscle Actuation Layer in isolation. (a,c) Head trajectories of the snake under two different initial HCO phases. (b,d) Corresponding membrane potentials (raster-style) for the four uncoupled HCOs. The plant already moves, but the individual patterns still need to be orchestrated.

7.3 Coordination Layer

Goal. Impose a forward/backward travelling wave across joints by (i) generating ordered spikes and (ii) converting them into HCO entrainment events.

Event level. At this level, we want to regulate the sequence of HCO activations. That is, beyond

left/right activation, we want to regulate ‘which joint should fire next’. We represent this as a traveling wave of spikes across joints. Each spike is an event that commands one joint to contract on one side. We also use a single burst event to tune the overall cycle period.

Functional level. The Coordination Layer is based on two main blocks:

Forward/Backward Generation Block (Figure 14). We build two 5-node RWTA rings: a *forward ring* and a *backward ring*. Each ring has four spiking neurons (one per joint) and one bursting neuron (which sets the dwell time and thus the gait frequency). All ten neurons (both rings) share a fast all-to-all inhibitory hub. This guarantees that only one neuron across *both* rings can be active at any given time (single-winner constraint). The forward ring produces a spike sequence in joint order $1 \rightarrow 2 \rightarrow 3 \rightarrow 4$; the backward ring produces $4 \rightarrow 3 \rightarrow 2 \rightarrow 1$.




Figure 14: Coordination Layer: Forward/Backward Generation Block (two RWTA 5-rings with shared inhibition) and the Interface Module Block (RF/LF chains plus AND-like selectors). The active ring sets the order of joints; the interface selects which side of each joint fires.

Interface Module Block (Figure 15). The spike events produced by the forward/backward generation block need to establish a consistent entrainment of the HCOs of the Muscle Actuation Layer. For each joint k , this is achieved by the interface module by the generation of two ‘entrainment events’: the spiking neuron RF_k inhibits the left HCO neuron L_k and excites the right HCO neuron R_k . The effect is to pull the joint k to the right. Similarly, the spiking neuron LF_k inhibits R_k and excites L_k , pulling the joint to the left.

The Interface Module decides to activate either the RF chain (RF_1 – RF_4) and the LF chain (LF_1 – LF_4). The activation is then propagated along the chain through the action of slow excitatory synapses. The activation decision is based on (i) the spike of the active ring of the forward/backward generation block, and on (ii) the associated HCO in the muscle actuation layer. For instance. The decision is implemented through four additional neurons, RFF , LFF , RFB , and LFB , acting as AND gates. The overall behavior of the network depends on the particular interconnection between modules, outlined in system level part.




Figure 15: Interface Module Block. RF_k nodes bias joint k towards “right” (inhibit left HCO neuron, excite right). LF_k nodes bias joint k towards “left”. Slow recurrent excitation within each chain ($RF_1 \leftrightarrow RF_2 \leftrightarrow RF_3 \leftrightarrow RF_4$, similarly LF) maintains the same entrainment type for the duration of a cycle once chosen by the first spike.

System level. The full coordination layer is obtained by interconnecting (i) the forward/backward generation block, (ii) the interface module block, and (iii) the muscle actuation layer block, as shown in Figure 16. The forward/backward generation block provides an ordered spike sequence (forward or backward wave), the interface block converts that sequence into joint-specific left/right entrainment events, and these entrainment events directly bias each HCO of the muscle actuation block, so that all four joints follow a coherent pattern. This conceptual picture is well represented by the top-to-bottom (excitatory and inhibitory) interconnections of Figure 16.

The returning connections between the output of the muscle actuation layer block and the input of the interface module block are used by the interface module block to decide the activation of the RF chain (RF1–RF4) and the LF chain (LF1–LF4). For instance, in forward mode (when the forward ring of the forward/backward generation block is active), the spike of neuron F1 is combined (AND operation) to the spike of the HCO neurons R1/L1, related to joint 1. If R1 is more depolarized than L1 when F1 spikes, the RF chain is enabled (right-bias); otherwise the LF chain is enabled (left-bias). In a similar way, in backward mode (when the backward ring is active of the forward/backward generation block is active), the same decision involves neuron B4 and the HCO neurons R4/L4 of joint 4.

Figure 17a illustrate the motion of the snake head when the coordination layer is active. The corresponding sequence of membrane potentials of the motor activation layer block neurons is provided in Figure 17b. At time $t = 45$ s an external reset is applied to the HCOs: R1–R4 are inhibited and L1–L4 activated. This demonstrates how the sequence of muscle activations is on one hand globally coordinated and on the other hand can be re-synchronized by a single event.




Figure 16: Full Coordination Layer interconnection. The Forward/Backward Generation Block provides ordered spikes. The Interface Module selects RF vs LF from the first spike and first HCO phase. The resulting entrainment events bias the HCOs in the Muscle Actuation Layer.




Figure 17: Effect of the Coordination Layer. A traveling wave of joint activations produces coherent locomotion, and a single external event can re-phase the whole body. a) Snake head trajectory and b) Membrane potentials of the four HCOs with Coordination Layer enabled. At $t = 45$ s an external pulse resets the HCO phases, switching the selected entrainment chain.

7.4 Supervisory Control Layer

Goal. When bumping against an obstacle, execute *right turn* or *back retreat*. Alternate between the two in case of multiple consecutive obstacles.

Event level. At the task level, the system must choose between two distinct behaviors upon encountering an obstacle. Each behavior is encoded as a persistent ON event generated downstream of the spiking network. This is achieved by passing the spiking output of an HCO through an ultra-slow excitatory synapse E_{us} . This synapse effectively functions as a low-pass filter, transforming the transient spikes into a sustained, step-like bias that signals the chosen behavior.

Functional level. The first component of the supervisory control layer is given by the **Switch block**, whose function is to store which behaviour (right turn vs back retreat) is currently active. As shown in Figure 18a, the switch block combines a RWTA network with ultra-slow synapses E_{us} acting on two new interface neurons, RTS (right-turn) and BRS (back-retreat).


The network's behavior is illustrated in Figure 18b. Activation of either the RT1-RT2 or BR1-BR2 pathway results in an ultra-slow output that biases the corresponding interface neuron, RTS or BRS, respectively. This bias primes the neuron to respond. For example, as shown in Figure 18b, even with identical inputs, the resulting activation levels of the RTS and BRS neurons are differentiated based on the current state of the RWTA network.

The second component of the supervisory control layer is the **falling-edge relay** represented by the dashed box in Figure 19. Its function is to detect *contact release*. When contact ends, each relay emits a rebound spike. This can be used to toggle the switch block, as outlined in the system level part below. The result is the alternation: first obstacle $\rightarrow$ right turn, next obstacle $\rightarrow$ back retreat, next obstacle $\rightarrow$ right turn again, and so on.

System level. The supervisory control layer interact with both the coordination layer and the muscle activation layer, as shown in Figure 19. The resulting behavior is described below. When contact is detected,

- (i) if neurons RT1 and RT2 are active, the RTS neuron becomes active. This neuron is directly connected to the muscles of the right-hand side of the snake via ultra-slow excitatory synapses E_{us} . The result is an ‘in-place’ right turn (Figure 19);
- (ii) if neurons BR1 and BR2 are active, the BRS neuron also becomes active. Its slow excitatory synapse triggers neuron BR4 of the coordination layer, causing a gait change. An extra self inhibition is also added to **BRS** to avoid repeated triggering of **B4**.

When contact ends, the “falling-edge relay” spikes and toggles the switch, so that the next contact will trigger the other behaviour. This is illustrated in Figure 20. The first obstacle triggers a right turn; after release, the internal state flips; the second obstacle triggers a back retreat.



(a) Two-state switch block. Ultra-slow excitation ($E_{u.s.}$) from the currently active state biases either RTS (right-turn strategy armed) or BRS (back-retreat strategy armed).

(b) Switch timing. The same sensor input is applied twice, but between them a spike toggles the internal state, so different strategies are expressed.

Figure 18: Supervisory Control Layer: the switch block maintains which avoidance strategy is “armed” and gates RTS/BRS accordingly.




Figure 19: System-level view of the Supervisory Control Layer. The tactile sensor excites RTS/BRS. ultra-slow excitatory synapses from RTS directly biases the muscle on the right hand side for an right turn. A slow excitatory synapses from BRS trigger the Coordination Layer (Backward ring) for a back retreat. The “falling-edge relay” neurons toggle the two-state switch on contact release, forcing alternation between strategies.




Figure 20: Behaviour under the Supervisory Control Layer. First obstacle contact: right turn. Contact ends, relay toggles internal state. Second obstacle contact: back retreat.

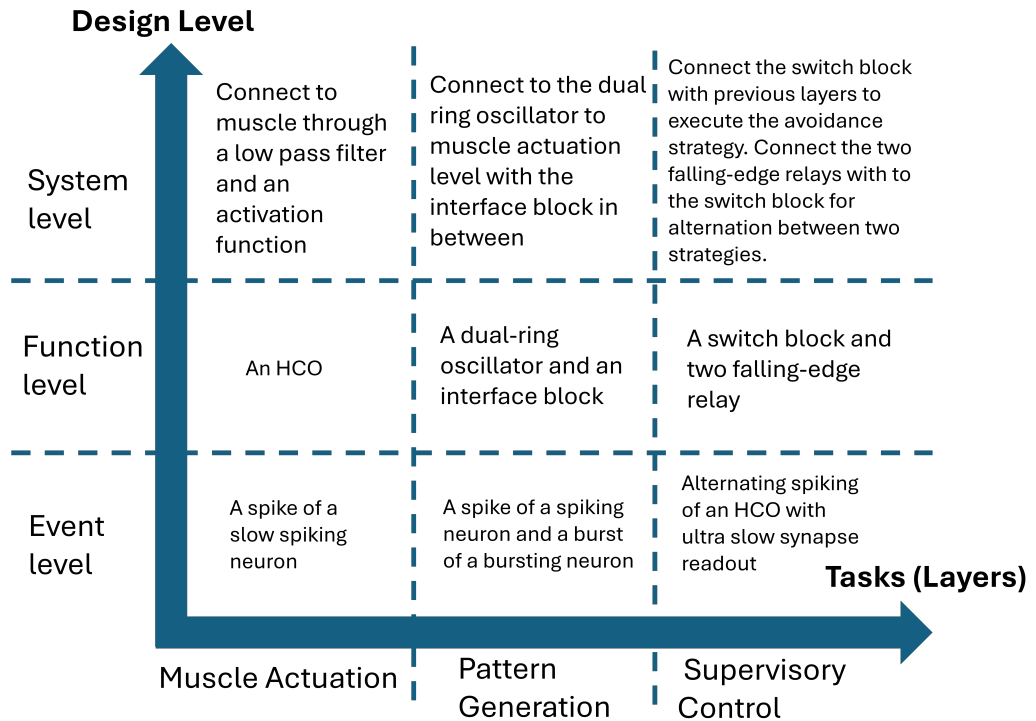


Figure 21: Summary of the controller architecture. The same RWTA design principles appear at three levels: Muscle Actuation Layer (local alternation), Coordination Layer (travelling wave and entrainment), Supervisory Control Layer (behavioural switching). All signals are voltage events (spikes or bursts); all interconnections are synaptic currents.

8 Discussion

Figure 21 can be regarded as a revision of the traditional design layout in Figure 1 in the event-based language of the present paper. A key feature of the rebound Winner-Take-All architecture is that all the cells in Figure 21 use the same modeling language and the same set of interconnection rules. Design, tuning, and adaptation are addressed at any level and any layer with the same concept of synaptically interconnected rebound neurons.

One methodology across layers and levels At the *event level*, all signals (sensory, internal, motor) are represented as spike or burst events generated by rebound-excitable neurons. At the *functional level*, the same rebound Winner-Take-All architecture implements both the rhythmic functionalities of pattern generation and the logical gates of an automaton. At the *system level*, sub-networks become states of a higher level automaton obeying the same architecture.

The unified language is in contrast to the traditional separation between the discrete language of automata for “higher level” design and the continuous (physical) language of differential equations for “lower level” design. Decision-making and regulation are merged in the same event-based framework, allowing continuous tuning and discrete topological/functional design at each scale.

Unified tuning. Because all layers and levels share the same substrate, they share the same tuning properties: (i) Changing neuron internal parameters to tune the shape and duration of an event (spike or burst) (ii) common or neuron-specific bias currents, which continuously tune the rebound time and thus rhythm frequency (demonstrated in Figure 9); (iii) transient sensory pulses, which apply immediate phase selection or hard resets by temporally biasing the current WTA competition. These three mechanisms are sufficient to modulate behaviours at *every* layer without introducing layer-specific control laws. As a result, adaptation is uniform: the same type of parameter change (bias or synaptic time constant/gain) has predictable, monotone effects from lower layer (muscle actuation) to higher layer (supervisory control).

How this differs from existing architectures. Conventional neuromorphic locomotion task typically couple (i) a continuous CPG for rhythms, (ii) a discrete finite-state machine for mode logic, and (iii) separate low-level joint controllers, glued together with cross-domain interfaces [32, 36, 37]. This heterogeneity forces different design languages (ODEs for regulation, automata or associative memories for logic [41, 42]), incurs interface fragility, and complicates tuning (parameters have layer-specific meanings). In contrast, the proposed framework avoids the need for ad-hoc glue logic between rate-based and event-based blocks, and no change of modeling paradigm between decision making and trajectory planning.

9 Conclusion

This paper has introduced *rebound winner-takes-all* networks as a unifying primitive for decision making and control. The neuromorphic architecture was illustrated on the design of a snake motion task. By designing the topology, we created a range of rebound winner-takes-all modules that deliver functions at different layers: state network, pattern generation and muscle actuation without the need for extra ingredients. And these three layers combine to form a complete controller for the snake motion task.

The framework opens a promising route toward large-scale neuromorphic event-based architectures. Because all computations are carried through physical voltages and currents variables, sensory signals can be injected *directly* into the network and motor commands can be read out *directly* from synaptic waveforms, eliminating the need for “recording and processing” and hence make the computation truly embodied and event-based.

From a controller point of view, the tunability of the network is not affected by the complexity of the topology—It is fully decentralized and distributed. This makes it appealing for event-based control as adapting the action time to the event time is completely decoupled from the orchestration task of the high level design.

Rebound Winner-Take-All networks sit at the *intersection* of continuous (physical) and discrete (algorithmic) paradigms. Winner-Take-All competition, long used as an analogue current-mode maximum selector, endows the network with logical decision capabilities, while continuous rebound dynamics enable tuneable rhythms unavailable to purely digital logic.

There are many questions that require research beyond the material of this paper. Primarily and foremost, we wish to develop learning and adaptation algorithms that can take advantage of the proposed neuromorphic architecture. The event-based nature of the neural network suggests the potential of Hebbian rules for learning and adaptation. Another key question is the hardware realisation of the “nervous system” architecture., and the respective advantages of digital and analog devices all the way from sensing to actuation. Those questions will be addressed in a future publication.

10 Appendix

10.1 Parameters for simulations and implementation detail for snake controller

All spiking neurons used in this paper has the following parameters: $\tau_s = 20$, $\alpha_f^- = \alpha_s^+ = 2$, $R = 0.5$, $\bar{V} = -1.5$. For bursting neuron in Figure 4, we use: $\tau_s = 5$, $\tau_{us} = 100$, $\alpha_f^- = \alpha_s^+ = 2$, $\alpha_s^- = 1.5$, $R = 0.5$, $\bar{V}_f = 3$, $\bar{V}_s = -1.5$. We use $\alpha_{us}^+ = 1.5$ for trajectory A and $\alpha_{us}^+ = 1.7$ for trajectory B.

In Figure 14, for HCO and the case of single inhibitory synapse the parameters are chosen to be $\tau = 0.5$, $\bar{V} = -1$ and $\alpha = 50$ for inhibitory synapses with weight -2 . All neurons are given a constant external input of -1.5 . In Figure 14, for other plots, all inhibitory synapse used the same parameters: $\tau = 0.5$, $\bar{V} = -2.5$ and $\alpha = 50$ with weight -2 . All excitatory synapses used $\tau = 20$, $\bar{V} = -2.5$ and $\alpha = 50$. All neurons are given a constant external input of -1.5 .

In Figure 10, all fast inhibitory synapse (blue lines) used the same parameters: $\tau = 0.5$, $\bar{V} = -2.5$ and $\alpha = 50$ with weight -2 . All slow inhibitory synapses (purple lines) used $\tau = 10$, $\bar{V} = -3$ and $\alpha = 50$ and synaptic weight -2.5 . All slow excitatory synapses (red arrow) used $\tau = 20$, $\bar{V} = -2.5$ and $\alpha = 50$ and synaptic weight 0.75 . All ultra-slow inhibitory and excitatory synapse (black lines and arrows) used the same parameters: $\tau = 200$, $\bar{V} = -4$ and $\alpha = 50$ with synaptic weight -2 for inhibition and synaptic weight 0.5 for excitation. All neurons are given a constant external input of -1.5 .

All the neurons and synapses are slowdown-ed by $3/100$ times in the snake simulation.

Muscle Actuation Layer In this layer, the bursting neurons have the following parameters: $\tau_s = 5$, $\tau_{us} = 100$, $\alpha_f^- = \alpha_s^+ = 2$, $\alpha_s^- = \alpha_{us}^+ = 1.5$, $R = 0.5$, $\bar{V}_f = 3$, $\bar{V}_s = -1.5$. Unless specified, the constant input for bursting neurons in this layer is $u = -1$. Each HCO use the I synapse as inhibition synapse with synaptic weight of -2 . The input of the muscle is directly connected to a synapse with $\tau = 20$, $\alpha = 16$, $\bar{V} = -1.875$ and gain 1.5 .

Coordination Layer In this layer, the bursting neurons have the following parameters: $\tau_s = 5$, $\tau_{us} = 100$, $\alpha_f^- = \alpha_s^+ = 2$, $\alpha_s^- = 1.1$, $\alpha_{us}^+ = 1.7$, $R = 0.5$, $\bar{V}_f = 3$, $\bar{V}_s = -1.5$. For neurons in the ring oscillators (both spiking and bursting neuron), we use constant input -0.5 . Fast all-to-all inhibition across the two rings are I synapses with synaptic weight of -2.5 . The slow excitatory connections in the two rings are E synapses with synaptic weight of 0.75 .

Neurons in the interface module are given constant input of -2 . The slow excitatory synapse (denote as red arrow in Figure 15) are E synapse with synaptic weight 1.5 . The fast excitatory synapse (denote as blue arrow in Figure 15) are E_{fl} synapse with synaptic weight 1.5 .

Figure 16 shows the connections between the ring oscillators, the interface module and the muscle actuation layer. The red arrows are slow excitatory connections E with synaptic weight 1 . The blue arrows from F1, R1, L1, B4, R4, L4 to RFF, LFF, RFB, LFB are E_f synapses. Specifically, the connections from F1, B4 to RFF, LFF, RFB, LFB have synaptic weight 1 . The connections from R1, L1, R4, L4 to RFF, LFF, RFB, LFB have synaptic weight of 0.6 . The blue arrows from RF1, LF1, RF2, LF2, RF3, LF3, RF4, LF4, are E_{fl} synapses with synaptic weight 1.5 . The blue circle head arrows from RF1, LF1, RF2, LF2, RF3, LF3, RF4, LF4, are I synapses with synaptic weight -3 .

Supervisory Control Layer RT1, RT2, BR1, BR2, BRR, RTR are given constant input of -0.5 . BRS and RTS are given constant input of -2 . Neurons in the 2-state machine (RT1, RT2, BR1, BR2) are all-to-all inhibited with I synapses with synaptic weight of -2.5 . The slow excitatory connection between them are denoted by red arrows in Figure 19 with synaptic weight of 0.75 . The sensor input to BRS and RTS is either ON (1) or off (0). The slow excitatory connection (E) between BRS and B4 has synaptic weight of 2 . The slow self inhibition (I_s) in BRS has synaptic weight -1.5 . The slow inhibition (I_s) from BRS, RTS to BRR, RTR have synaptic weight of -1 . The slow excitatory connection (E) from BRR, RTR to RT1, BR1 have synaptic weight of 2 . The ultra-slow activation (E_{us}) from RT2, BR2 to BRS, RTS have synaptic weight of 1 . The ultra-slow connection between RTS and the muscles have $\tau = 100$, $\alpha = 16$, $\bar{V} = 2.625$ and synaptic weight of 0.5 .

For Figure 20, the α_{us}^+ of the bursting neurons in the muscle actuation layer is adjusted to 2.75 .

References

- [1] R. Sepulchre, G. Drion, and A. Franci. “Control Across Scales by Positive and Negative Feedback”. en. In: *Annual Review of Control, Robotics, and Autonomous Systems* 2.1 (2019), pp. 89–113. ISSN: 2573-5144, 2573-5144. DOI: [10.1146/annurev-control-053018-023708](https://doi.org/10.1146/annurev-control-053018-023708).
- [2] N. Matni, A. D. Ames, and J. C. Doyle. “A Quantitative Framework for Layered Multirate Control: Toward a Theory of Control Architecture”. In: *IEEE Control Systems Magazine* 44.3 (2024), pp. 52–94. ISSN: 1941-000X. DOI: [10.1109/MCS.2024.3382388](https://doi.org/10.1109/MCS.2024.3382388).
- [3] R. Sepulchre. “Spiking Control Systems”. In: *Proceedings of the IEEE* 110.5 (2022), pp. 577–589. ISSN: 1558-2256. DOI: [10.1109/JPROC.2022.3163926](https://doi.org/10.1109/JPROC.2022.3163926).

- [4] L. Ribar and R. Sepulchre. “Neuromorphic Control: Designing Multiscale Mixed-Feedback Systems”. In: *IEEE Control Systems Magazine* 41.6 (2021), pp. 34–63. ISSN: 1941-000X. DOI: [10.1109/MCS.2021.3107560](https://doi.org/10.1109/MCS.2021.3107560).
- [5] R. Schmetterling, F. Forni, A. Franci, and R. Sepulchre. “Neuromorphic Control of a Pendulum”. In: *IEEE Control Systems Letters* 8 (2024), pp. 1235–1240. ISSN: 2475-1456. DOI: [10.1109/LCSYS.2024.3409093](https://doi.org/10.1109/LCSYS.2024.3409093).
- [6] C. Mead. *Analog VLSI and neural systems*. eng. Open Library ID: OL2038266M. Reading, Mass: Addison-Wesley, 1989. ISBN: 9780201059922.
- [7] G. Indiveri. “Winner-Take-All Networks with Lateral Excitation”. en. In: *Analog Integrated Circuits and Signal Processing* 13.1 (1997), pp. 185–193. ISSN: 1573-1979. DOI: [10.1023/A:1008200515504](https://doi.org/10.1023/A:1008200515504).
- [8] G. Indiveri et al. “Neuromorphic Silicon Neuron Circuits”. English. In: *Frontiers in Neuroscience* 5 (2011). ISSN: 1662-453X. DOI: [10.3389/fnins.2011.00073](https://doi.org/10.3389/fnins.2011.00073).
- [9] B. V. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. R. Chandrasekaran, J.-M. Bussat, R. Alvarez-Icaza, J. V. Arthur, P. A. Merolla, and K. Boahen. “Neurogrid: A Mixed-Analog-Digital Multichip System for Large-Scale Neural Simulations”. In: *Proceedings of the IEEE* 102.5 (2014), pp. 699–716. ISSN: 1558-2256. DOI: [10.1109/JPROC.2014.2313565](https://doi.org/10.1109/JPROC.2014.2313565).
- [10] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana. “The SpiNNaker Project”. In: *Proceedings of the IEEE* 102.5 (2014), pp. 652–665. ISSN: 1558-2256. DOI: [10.1109/JPROC.2014.2304638](https://doi.org/10.1109/JPROC.2014.2304638).
- [11] D. Kudithipudi et al. “Neuromorphic computing at scale”. en. In: *Nature* 637.8047 (2025), pp. 801–812. ISSN: 0028-0836, 1476-4687. DOI: [10.1038/s41586-024-08253-8](https://doi.org/10.1038/s41586-024-08253-8).
- [12] T. G. Brown. “The intrinsic factors in the act of progression in the mammal”. en. In: *Proceedings of the Royal Society of London. Series B, Containing Papers of a Biological Character* 84.572 (1911), pp. 308–319. ISSN: 0950-1193, 2053-9185. DOI: [10.1098/rspb.1911.0077](https://doi.org/10.1098/rspb.1911.0077).
- [13] A. J. Ijspeert. “Central pattern generators for locomotion control in animals and robots: A review”. In: *Neural Networks. Robotics and Neuroscience* 21.4 (2008), pp. 642–653. ISSN: 0893-6080. DOI: [10.1016/j.neunet.2008.03.014](https://doi.org/10.1016/j.neunet.2008.03.014).
- [14] A. J. Ijspeert, A. Crespi, and J.-M. Cabelguen. “Simulation and robotics studies of salamander locomotion”. en. In: *Neuroinformatics* 3.3 (2005), pp. 171–195. ISSN: 1559-0089. DOI: [10.1385/NI:3:3:171](https://doi.org/10.1385/NI:3:3:171).
- [15] K. Matsuoka. “Sustained oscillations generated by mutually inhibiting neurons with adaptation”. en. In: *Biological Cybernetics* 52.6 (1985), pp. 367–376. ISSN: 1432-0770. DOI: [10.1007/BF00449593](https://doi.org/10.1007/BF00449593).
- [16] K. Matsuoka. “Analysis of a neural oscillator”. eng. In: *Biological Cybernetics* 104.4-5 (2011), pp. 297–304. ISSN: 1432-0770. DOI: [10.1007/s00422-011-0432-z](https://doi.org/10.1007/s00422-011-0432-z).
- [17] W. Maass. “On the Computational Power of Winner-Take-All”. In: *Neural Computation* 12.11 (2000), pp. 2519–2535. ISSN: 0899-7667. DOI: [10.1162/089976600300014827](https://doi.org/10.1162/089976600300014827).
- [18] M. Usher and J. L. McClelland. “The time course of perceptual choice: the leaky, competing accumulator model”. eng. In: *Psychological Review* 108.3 (2001), pp. 550–592. ISSN: 0033-295X. DOI: [10.1037/0033-295x.108.3.550](https://doi.org/10.1037/0033-295x.108.3.550).
- [19] J. Lazzaro, S. Ryckebusch, M. Mahowald, and C. A. Mead. “Winner-Take-All Networks of O(N) Complexity”. In: *Advances in Neural Information Processing Systems*. Vol. 1. Morgan-Kaufmann, 1988.
- [20] E. Neftci, J. Binas, U. Rutishauser, E. Chicca, G. Indiveri, and R. J. Douglas. “Synthesizing cognition in neuromorphic electronic systems”. en. In: *Proceedings of the National Academy of Sciences* 110.37 (2013). ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.1212083110](https://doi.org/10.1073/pnas.1212083110).
- [21] S. Grossberg. “Adaptive pattern classification and universal recoding: I. Parallel development and coding of neural feature detectors”. en. In: *Biological Cybernetics* 23.3 (1976), pp. 121–134. ISSN: 1432-0770. DOI: [10.1007/BF00344744](https://doi.org/10.1007/BF00344744).
- [22] T. G. Brown. “On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of rhythmic activity in progression, and a theory of the evolution of function in the nervous system”. en. In: *The Journal of Physiology* 48.1 (1914), pp. 18–46. ISSN: 0022-3751, 1469-7793. DOI: [10.1113/jphysiol.1914.sp001646](https://doi.org/10.1113/jphysiol.1914.sp001646).

- [23] D. M. Wilson. “The Central Nervous Control of Flight in A Locust”. en. In: *Journal of Experimental Biology* 38.2 (1961), pp. 471–490. ISSN: 0022-0949, 1477-9145. DOI: [10.1242/jeb.38.2.471](https://doi.org/10.1242/jeb.38.2.471).
- [24] S. Rossignol, R. Dubuc, and J.-P. Gossard. “Dynamic sensorimotor interactions in locomotion”. eng. In: *Physiological Reviews* 86.1 (2006), pp. 89–154. ISSN: 0031-9333. DOI: [10.1152/physrev.00028.2005](https://doi.org/10.1152/physrev.00028.2005).
- [25] P. A. Guertin. “The mammalian central pattern generator for locomotion”. eng. In: *Brain Research Reviews* 62.1 (2009), pp. 45–56. ISSN: 1872-6321. DOI: [10.1016/j.brainresrev.2009.08.002](https://doi.org/10.1016/j.brainresrev.2009.08.002).
- [26] E. Marder and D. Bucher. “Central pattern generators and the control of rhythmic movements”. In: *Current Biology* 11.23 (2001), R986–R996. ISSN: 0960-9822. DOI: [10.1016/S0960-9822\(01\)00581-4](https://doi.org/10.1016/S0960-9822(01)00581-4).
- [27] B. Dirk, H. Gal, G. Jorge, and N. Farzan. “Central Pattern Generators”. In: *Encyclopedia of Life Sciences* (2015), pp. 1–12. DOI: [10.1002/9780470015902.a0000032.pub2](https://doi.org/10.1002/9780470015902.a0000032.pub2).
- [28] F. K. Skinner, N. Kopell, and E. Marder. “Mechanisms for oscillation and frequency control in reciprocally inhibitory model neural networks”. en. In: *Journal of Computational Neuroscience* 1.1 (1994), pp. 69–87. ISSN: 1573-6873. DOI: [10.1007/BF00962719](https://doi.org/10.1007/BF00962719).
- [29] E. Morozova, P. Newstein, and E. Marder. “Reciprocally inhibitory circuits operating with distinct mechanisms are differently robust to perturbation and modulation”. In: *eLife* 11 (2022). Ed. by R. L. Calabrese, R. L. Calabrese, P. Katz, and J. Marino Ramirez. Publisher: eLife Sciences Publications, Ltd, e74363. ISSN: 2050-084X. DOI: [10.7554/eLife.74363](https://doi.org/10.7554/eLife.74363).
- [30] M. Lewis, R. Etienne-Cummings, A. Cohen, and M. Hartmann. “Toward biomorphic control using custom aVLSI CPG chips”. In: *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*. Vol. 1. ISSN: 1050-4729. 2000, 494–500 vol.1. DOI: [10.1109/ROBOT.2000.844103](https://doi.org/10.1109/ROBOT.2000.844103).
- [31] J. Yu, M. Tan, J. Chen, and J. Zhang. “A Survey on CPG-Inspired Control Models and System Implementation”. In: *IEEE Transactions on Neural Networks and Learning Systems* 25.3 (2014), pp. 441–456. ISSN: 2162-2388. DOI: [10.1109/TNNLS.2013.2280596](https://doi.org/10.1109/TNNLS.2013.2280596).
- [32] Á. Lozano, M. Rodríguez, and R. Barrio. “Control strategies of 3-cell Central Pattern Generator via global stimuli”. en. In: *Scientific Reports* 6.1 (2016). Publisher: Nature Publishing Group, p. 23622. ISSN: 2045-2322. DOI: [10.1038/srep23622](https://doi.org/10.1038/srep23622).
- [33] L. Luo, Z. Wei, Y. Sun, W. Zhang, and J. Zhou. “Preliminary Design and Control of an Ankle Rehabilitation Robot With Variable Stiffness Actuator”. In: *IEEE Access* 12 (2024), pp. 195906–195917. ISSN: 2169-3536. DOI: [10.1109/ACCESS.2024.3517746](https://doi.org/10.1109/ACCESS.2024.3517746).
- [34] P. López-Osorio, J. P. Domínguez-Morales, and F. Perez-Peña. “A Neuromorphic Vision and Feedback Sensor Fusion Based on Spiking Neural Networks for Real-Time Robot Adaption”. en. In: *Advanced Intelligent Systems* 6.5 (2024), p. 2300646. ISSN: 2640-4567, 2640-4567. DOI: [10.1002/aisy.202300646](https://doi.org/10.1002/aisy.202300646).
- [35] I. Abadía, A. Bruel, G. Courtine, A. J. Ijspeert, E. Ros, and N. R. Luque. “A neuromechanics solution for adjustable robot compliance and accuracy”. en. In: *Science Robotics* 10.98 (2025), eadp2356. ISSN: 2470-9476. DOI: [10.1126/scirobotics.adp2356](https://doi.org/10.1126/scirobotics.adp2356).
- [36] Z. Zhang, G. Bellegarda, M. Shafiee, and A. Ijspeert. “Online Optimization of Central Pattern Generators for Quadruped Locomotion”. In: *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. ISSN: 2153-0866. 2024, pp. 13547–13554. DOI: [10.1109/IROS58592.2024.10802762](https://doi.org/10.1109/IROS58592.2024.10802762).
- [37] A. Linares-Barranco, E. Pinero-Fuentes, S. Canas-Moreno, A. Rios-Navarro, Maryada, C. Wu, J. Zhao, D. Zendrikov, and G. Indiveri. *Towards hardware Implementation of WTA for CPG-based control of a Spiking Robotic Arm*. arXiv:2202.07064. 2022. DOI: [10.48550/arXiv.2202.07064](https://doi.org/10.48550/arXiv.2202.07064).
- [38] E. Majani, R. Erlanson, and Y. Abu-Mostafa. “On the K-Winners-Take-All Network”. In: *Advances in Neural Information Processing Systems*. Vol. 1. Morgan-Kaufmann, 1988.
- [39] M. Oster, R. Douglas, and S.-C. Liu. “Computation with Spikes in a Winner-Take-All Network”. In: *Neural Computation* 21.9 (2009), pp. 2437–2465. ISSN: 0899-7667. DOI: [10.1162/neco.2009.07-08-829](https://doi.org/10.1162/neco.2009.07-08-829).

- [40] G. Indiveri, E. Chicca, and R. J. Douglas. “Artificial Cognitive Systems: From VLSI Networks of Spiking Neurons to Neuromorphic Cognition”. en. In: *Cognitive Computation* 1.2 (2009), pp. 119–127. ISSN: 1866-9964. DOI: [10.1007/s12559-008-9003-6](https://doi.org/10.1007/s12559-008-9003-6).
- [41] M. Cotteret, H. Grotto, M. Ziegler, and E. Chicca. “Vector Symbolic Finite State Machines in Attractor Neural Networks”. en. In: *Neural Computation* 36.4 (2024), pp. 549–595. ISSN: 0899-7667, 1530-888X. DOI: [10.1162/neco_a_01638](https://doi.org/10.1162/neco_a_01638).
- [42] M. Cotteret, H. Grotto, A. Renner, J. Chen, E. Neftci, H. Wu, G. Indiveri, M. Ziegler, and E. Chicca. “Distributed representations enable robust multi-timescale symbolic computation in neuromorphic hardware”. In: *Neuromorphic Computing and Engineering* 5.1 (2025), p. 014008. ISSN: 2634-4386. DOI: [10.1088/2634-4386/ada851](https://doi.org/10.1088/2634-4386/ada851).
- [43] J. J. Hopfield. “Neural networks and physical systems with emergent collective computational abilities.” en. In: *Proceedings of the National Academy of Sciences* 79.8 (1982), pp. 2554–2558. ISSN: 0027-8424, 1091-6490. DOI: [10.1073/pnas.79.8.2554](https://doi.org/10.1073/pnas.79.8.2554).
- [44] J. Abrahamsen, P. Hafliger, and T. Lande. “A time domain winner-take-all network of integrate-and-fire neurons”. In: *2004 IEEE International Symposium on Circuits and Systems (IEEE Cat. No.04CH37512)*. Vol. 5. 2004, pp. V–V. DOI: [10.1109/ISCAS.2004.1329537](https://doi.org/10.1109/ISCAS.2004.1329537).
- [45] M. Rabinovich, A. Volkovskii, P. Lecanda, R. Huerta, H. D. I. Abarbanel, and G. Laurent. “Dynamical Encoding by Networks of Competing Neuron Groups: Winnerless Competition”. en. In: *Physical Review Letters* 87.6 (2001), p. 068102. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.87.068102](https://doi.org/10.1103/PhysRevLett.87.068102).
- [46] V. S. Afraimovich, M. I. Rabinovich, and P. Varona. “HETEROCLINIC CONTOURS IN NEURAL ENSEMBLES AND THE WINNERLESS COMPETITION PRINCIPLE”. en. In: *International Journal of Bifurcation and Chaos* 14.04 (2004), pp. 1195–1208. ISSN: 0218-1274, 1793-6551. DOI: [10.1142/S0218127404009806](https://doi.org/10.1142/S0218127404009806).
- [47] M. I. Rabinovich and P. Varona. “Discrete Sequential Information Coding: Heteroclinic Cognitive Dynamics”. English. In: *Frontiers in Computational Neuroscience* 12 (2018). ISSN: 1662-5188. DOI: [10.3389/fncom.2018.00073](https://doi.org/10.3389/fncom.2018.00073).
- [48] C. Bick and M. I. Rabinovich. “Dynamical Origin of the Effective Storage Capacity in the Brain’s Working Memory”. en. In: *Physical Review Letters* 103.21 (2009), p. 218101. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.103.218101](https://doi.org/10.1103/PhysRevLett.103.218101).
- [49] P. Seliger, L. S. Tsimring, and M. I. Rabinovich. “Dynamics-based sequential memory: Winnerless competition of patterns”. en. In: *Physical Review E* 67.1 (2003), p. 011905. ISSN: 1063-651X, 1095-3787. DOI: [10.1103/PhysRevE.67.011905](https://doi.org/10.1103/PhysRevE.67.011905).
- [50] P. Ashwin and C. Postlethwaite. “On designing heteroclinic networks from graphs”. In: *Physica D: Nonlinear Phenomena* 265 (2013), pp. 26–39. ISSN: 0167-2789. DOI: [10.1016/j.physd.2013.09.006](https://doi.org/10.1016/j.physd.2013.09.006).
- [51] M. I. Rabinovich, Y. Sokolov, and R. Kozma. “Robust sequential working memory recall in heterogeneous cognitive networks”. In: *Frontiers in Systems Neuroscience* 8 (2014), p. 220. ISSN: 1662-5137. DOI: [10.3389/fnsys.2014.00220](https://doi.org/10.3389/fnsys.2014.00220).
- [52] P. Ashwin, Ö. Karabacak, and T. Nowotny. “Criteria for robustness of heteroclinic cycles in neural microcircuits”. In: *The Journal of Mathematical Neuroscience* 1.1 (2011), p. 13. ISSN: 2190-8567. DOI: [10.1186/2190-8567-1-13](https://doi.org/10.1186/2190-8567-1-13).
- [53] L. Herron, P. Sartori, and B. Xue. “Robust Retrieval of Dynamic Sequences through Interaction Modulation”. en. In: *PRX Life* 1.2 (2023), p. 023012. ISSN: 2835-8279. DOI: [10.1103/PRXLife.1.023012](https://doi.org/10.1103/PRXLife.1.023012).
- [54] A. D. Horschler, K. A. Daltorio, H. J. Chiel, and R. D. Quinn. “Designing responsive pattern generators: stable heteroclinic channel cycles for modeling and control”. In: *Bioinspiration & Biomimetics* 10.2 (2015), p. 026001. ISSN: 1748-3190. DOI: [10.1088/1748-3190/10/2/026001](https://doi.org/10.1088/1748-3190/10/2/026001).
- [55] A. L. Hodgkin and A. F. Huxley. “A quantitative description of membrane current and its application to conduction and excitation in nerve”. In: *The Journal of Physiology* 117.4 (1952), pp. 500–544. ISSN: 0022-3751.
- [56] R. Sepulchre, G. Drion, and A. Franci. “Excitable Behaviors”. en. In: *Emerging Applications of Control and Systems Theory: A Festschrift in Honor of Mathukumalli Vidyasagar*. Ed. by R. Tempo, S. Yurkovich, and P. Misra. Cham: Springer International Publishing, 2018, pp. 269–280. ISBN: 9783319670683. DOI: [10.1007/978-3-319-67068-3_20](https://doi.org/10.1007/978-3-319-67068-3_20).

- [57] Y. Huo, F. Forni, and R. Sepulchre. *A Winner-Takes-All Mechanism for Event Generation*. arXiv:2504.11374. 2025. DOI: [10.48550/arXiv.2504.11374](https://doi.org/10.48550/arXiv.2504.11374).
- [58] H. R. Wilson and J. D. Cowan. “Excitatory and Inhibitory Interactions in Localized Populations of Model Neurons”. In: *Biophysical Journal* 12.1 (1972), pp. 1–24. ISSN: 0006-3495. DOI: [10.1016/S0006-3495\(72\)86068-5](https://doi.org/10.1016/S0006-3495(72)86068-5).
- [59] H. R. (R. Wilson. *Spikes, decisions, and actions : the dynamical foundations of neuroscience*. Oxford University Press, 1999.
- [60] G. Drion, J. Dethier, A. Franci, and R. Sepulchre. “Switchable slow cellular conductances determine robustness and tunability of network states”. en. In: *PLOS Computational Biology* 14.4 (2018), e1006125. ISSN: 1553-7358. DOI: [10.1371/journal.pcbi.1006125](https://doi.org/10.1371/journal.pcbi.1006125).
- [61] M. Mirbagheri, H. Barbeau, and R. Kearney. “Intrinsic and reflex contributions to human ankle stiffness: variation with activation level and position”. en. In: *Experimental Brain Research* 135.4 (2000), pp. 423–436. ISSN: 1432-1106. DOI: [10.1007/s002210000534](https://doi.org/10.1007/s002210000534).
- [62] A. G. Feldman. “Once more on the equilibrium-point hypothesis (lambda model) for motor control”. eng. In: *Journal of Motor Behavior* 18.1 (1986), pp. 17–54. ISSN: 0022-2895. DOI: [10.1080/00222895.1986.10735369](https://doi.org/10.1080/00222895.1986.10735369).