Rebound Modern Hopfield Network

Yongkang Huo [(1)] and Rodolphe Sepulchre [(2)]

[(1)] University of Cambridge, [(2)] KU Leuven

In the age of digital computing, robotics has become increasingly divided into two domains. High-level tasks, such as planning and decision-making, are typically addressed using discrete methodologies, whereas low-level tasks, such as regulation and tracking, are handled using continuous methodologies. This division separates automation from regulation, and decision-making from physical interaction. While such a separation offers advantages in analysis and design, it comes at the cost of an unnatural and often difficult translation layer between the two. This interface has become a bottleneck in the development of scalable architectures capable of both control and regulation across a wide range of spatial and temporal scales [1, 2]. We argue that there should be no hard boundary between the discrete and the continuous. Instead, we describe automation and regulation in the unified language of multi-timescale events, in which the classical problems of automation and regulation lie at opposite ends of a temporal continuum: continuous regulation corresponds to fast events, while discrete automation corresponds to slow events. Neuromorphic control [3–5] adopts this principle: it mixes 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].

Our previous work, A Neuromorphic Architecture for Scalable Event-Based Control [12], introduced the rebound Winner-Take-All (RWTA) motif as the basic building block of a scalable event-based control architecture. A Winner-Take-All (WTA) network is a mutually inhibitory network in which competing units suppress one another so that only one state, or one admissible group of states, remains active at a time, providing a reliable mechanism for discrete selection and state-machine-like computation. Rebound refers to post-inhibitory rebound excitability: a unit generates an event when an inhibitory input is released, rather than only in response to direct excitation. The property has long been studied in neurophysiology and has inspired the design of artificial central pattern generators in rhythmic machines [13–16]. By combining rebound excitability at the cellular level with WTA computation at the network level, the RWTA motif separates event generation from event orchestration, allowing rhythmic generation and decision-making to be described within the same event-based physical framework. This architecture was then organized hierarchically and illustrated through the nervous-system design of a snake robot. A central conclusion of that work is that the architecture separates event orchestration from event generation: the topology of the network, and more generally of the network-of-networks at slower time scales, determines the pattern and order of events, whereas the neuronal dynamics at faster time scales determine the shape and duration of each event. Because this separation repeats hierarchically across scales, higher-level events can be constructed as ordered sequences of lower-level events within a single physical modeling language.

The present work continues the investigation of this neuromorphic architecture by generalizing the winner-take-all topology to topologies that can encode arbitrary discrete patterns with a learning strategy inspired from the modern Hopfield model. Modern Hopfield network [17] is a two-layer associative memory consisting of a visible layer and a hidden layer. The visible layer represents the current sensory, motor, or internal state, while the hidden layer represents stored features, memories, or latent states selected by that visible activity. In common two-layer realizations, within-groups all-to-all inhibition can be introduced to hidden units to induce competition in the hidden layer, so that the architecture can be viewed as a collection of competing hidden subnetworks coupled through a shared visible substrate. From this perspective, the modern Hopfield network can be regarded as several Winner-Take-All networks that share a common visible layer. This viewpoint provides a direct architectural bridge from the rebound Winner-Take-All motif to a rebound modern Hopfield network. By replacing the neurons in the modern Hopfield network with rebound neurons, we recover a structure analogous to the Winner-Take-All case: rebound dynamics remain responsible for event generation, while the shared visible layer allows multiple hidden competitions to coordinate and encode richer spatial patterns than a single Winner-Take-All network can represent.

Learning can then be introduced at the level of spatial pattern formation by adapting the symmetric connections between visible and hidden neurons. In its simplest form, this can be achieved through Hebbian learning [18], and it can be generalized to equilibrium propagation [19] when a cost function is specified. The decoupling between event generation and event orchestration allows us to use traditional learning methods, such as Hebbian learning or equilibrium propagation, to learn the topology that defines the spatial pattern, while leaving the tuning of event shape and event timing to the excitable dynamics of the neurons and the timescales of the synapses.

Feedback can also be incorporated into the network in a more explicit way. As noted above, the visible layer represents the current sensory, motor, or internal state, which in turn determines the hidden state and can be linked to action. On top of this, we introduce slow excitatory connections between hidden neurons. These connections provide a feedforward pathway through the network, so that, in the absence of feedback, the network generates temporal patterns determined by this feedforward pathway. When feedback is present, however, the resulting pattern emerges from the competition between feedback strength and feedforward strength.

Overall, the rebound modern Hopfield network extends the original rebound Winner-Take-All framework from a hand-designed event architecture to a learnable and more expressive event-based control architecture. By preserving the separation between event generation and event orchestration, while enriching the spatial patterns that can be encoded and introducing explicit feedback, it provides a route toward scalable neuromorphic controllers that can integrate regulation, decision-making, memory, and action within a single framework.

[1] R. Sepulchre et al. en. In: *Annual Review of Control, Robotics, and Autonomous Systems* 2.1 (2019), pp. 89–113.

[2] N. Matni et al. In: *IEEE Control Systems Magazine* 44.3 (2024), pp. 52–94.

[3] R. Sepulchre. In: *Proceedings of the IEEE* 110.5 (2022), pp. 577–589.

[4] L. Ribar and R. Sepulchre. In: *IEEE Control Systems Magazine* 41.6 (2021), pp. 34–63.

[5] R. Schmetterling et al. In: *IEEE Control Systems Letters* 8 (2024), pp. 1235–1240.

[6] C. Mead. *Analog VLSI and neural systems*. eng. Open Library ID: OL2038266M. Reading, Mass: Addison-Wesley, 1989.

[7] G. Indiveri. en. In: *Analog Integrated Circuits and Signal Processing* 13.1 (1997), pp. 185–193.

[8] G. Indiveri et al. English. In: *Frontiers in Neuroscience* 5 (2011).

[9] B. V. Benjamin et al. In: *Proceedings of the IEEE* 102.5 (2014), pp. 699–716.

[10] S. B. Furber et al. In: *Proceedings of the IEEE* 102.5 (2014), pp. 652–665.

[11] D. Kudithipudi et al. en. In: *Nature* 637.8047 (2025), pp. 801–812.

[12] Y. Huo et al. In: *Neuromorphic Computing and Engineering* 6.1 (2026), p. 014018.

[13] A. J. Ijspeert. In: *Neural Networks*. Robotics and Neuroscience 21.4 (2008), pp. 642–653.

[14] A. J. Ijspeert et al. en. In: *Neuroinformatics* 3.3 (2005), pp. 171–195.

[15] K. Matsuoka. en. In: *Biological Cybernetics* 52.6 (1985), pp. 367–376.

[16] K. Matsuoka. eng. In: *Biological Cybernetics* 104.4-5 (2011), pp. 297–304.

[17] D. Krotov and J. J. Hopfield. In: *Advances in Neural Information Processing Systems*. Vol. 29. Curran Associates, Inc., 2016.

[18] D. Hebb. *The Organization of Behavior*. en. 0th ed. Psychology Press, 2005.

[19] B. Scellier and Y. Bengio. English. In: *Frontiers in Computational Neuroscience* 11 (2017).