

Spatiotemporal Spike Pattern Detection with Second-order Memristive Synapses

Yuting Wu, Sangmin Yoo, Fan-hsuan Meng, and Wei D. Lu*
 Department of Electrical Engineering and Computer Science
 The University of Michigan
 Ann Arbor, MI 48109, USA
 *E-mail: wluee@umich.edu

Abstract—Spatiotemporal patterns of spike trains convey critical information in a biological neural network. Second-order memristive devices, whose internal state variables offer short- and long-term temporal dynamics, have been employed to natively decode the temporal correlation of spiking patterns through bio-realistic implementation of synaptic learning rules. In this work, we demonstrate that a single artificial post-synaptic neuron equipped with an array of second-order memristive synapses can localize a precise spatiotemporal firing pattern, which repeats irregularly within an equally dense background of Poisson spiking events, in an unsupervised fashion.

Keywords—Second-order memristor, STDP, SNN, unsupervised learning

I. INTRODUCTION

In biological networks, sparse, binary spiking events carry critical information and induce synaptic plasticity, allowing the system to achieve remarkable processing and learning capabilities while consuming little energy. Taking inspiration from biology, spiking neural networks (SNNs) are promising for the efficient implementation of neural networks. However, synaptic learning rules, such as spike-timing dependent plasticity (STDP), typically requires complex circuit design to implement and induces significant overhead in area and power consumption. Second-order memristive devices[1], where the evolution of the second-order state variable natively possesses the required temporal dynamics, can bio-realistically emulate synaptic plasticity with simple pulses, achieving synaptic functions in a compact and efficient manner. Memristor-based networks with STDP function have been employed for temporal information processing, including correlation pattern detection[2, 3], synaptic connection pattern reconstruction[4, 5] and coincidence detection[6].

In biology, repeated occurrence of precise spatiotemporal patterns has been detected across cortical regions and is believed to encode critical information for downstream processing[7-9]. Yet, it is computationally expensive to extract the spiking patterns from a very noisy background. Surprisingly, it has been shown theoretically that a single post-synaptic neuron can learn through STDP rule to identify the repeated precise pattern with high fidelity and selectivity[10]. In this article, we aim to demonstrate how this information processing capability can be realized by a second-order memristor network in a bio-plausible fashion with device-model based simulation. The effects of device non-idealities on the detection accuracy will also be investigated.

II. SECOND-ORDER MEMRISTOR DEVICE AND MODEL

A. Second-order Memristor with STDP

STDP learning rule has been implemented in a number of studies with memristors[12-15]. However, precise pulse shape and overlap need to be engineered in a first-order device due to the lack of native mechanisms to encode timing

information. In a second-order memristor, the native temporal dynamics of a second state variable such as the local temperature can be used to natively implement STDP learning rule with simple, non-overlapping pulses, as shown in Fig. 1a. Additionally, the STDP time constants can be controlled through device structure or material engineering, for example, by inserting a heat insulation (HI) layer that can slow down the temperature decay process. Fig. 1b shows the measured STDP characteristics in such a tantalum-oxide based memristor device, fabricated with and without the inserted HI layer, respectively. The measured data can be well-fitted curves resembling those observed in biological




Fig. 1. (a) Top view of a fabricated second-order memristor device with heat insulation (HI) layers. (b) Device structure, showing the switching layer and the HI layers. (c) Temperature profile of the device with and without HI layers, simulated immediately after the forming pulse. (d) Pulse schemes with pre- and post-synaptic spikes. The spontaneous decay of the internal temperature naturally encodes the relative timing information, and the enhanced temperature together with the input spike collaboratively induce the conductance change. (e) Natively implemented STDP characteristics of memristors with (blue) and without (red) HI layers. Dots with error bar are experimental data and the curves are fitted STDP functions.

synapses. Fig. 1c shows the pre- and post-spikes used in the measurements, and the evolution of the second-order state variable (temperature) due to Joule heating and spontaneous relaxation, as well as the response of the first-order state variable (conductive filament size) that natively decodes the relative timing information in the pre- and post-spike pairs [1].

III. MEMRISTOR NETWORK FOR SPATIOTEMPORAL PATTERN DETECTION

A. System Overview

Fig. 2a illustrates the structure of the proposed memristor-network. The top electrode (TE) of each memristive synapses receives the spikes from one pre-synaptic neuron. The 1,000 pre-synaptic neurons fire spikes following Poisson distributions at similar rates of approximately 2,000 Hz for the entire 4.5s simulation duration. For a specific time period (i.e. 0.5ms), a selected group of the input neurons fire a deterministic spiking pattern instead, while the unselected neurons keep firing Poisson spikes. In the baseline test case,

the selected neurons correspond to 50% of the entire population and were chosen randomly. The precise firing pattern repeats with irregular intervals, and the average occurrence rate is 0.25 in the baseline case. Fig. 2b shows the raster plot of 100 randomly selected neurons, with neurons 0-49 involved in the pattern and neurons 50-99 not involved in the pattern. The target spatiotemporal pattern (marked by red)




Fig. 2. (a) Overview of the memristor network. A post-synaptic LIF neuron is connected with an array of second-order memristive synapses and receives incoming spikes from 1,000 afferents. (b) Raster plot of 100 random selected neurons, with 50 neurons in the selected group. The 0.5ms long firing pattern is repeated at irregular intervals with baseline probability of 0.25 throughout the entire 4.5s simulation duration. In the plot, the neurons have been re-ordered so that neurons involved in the pattern are shown as neurons 0-49, and the pattern is marked in red. The bottom plot shows the average firing rate of all afferents over a 0.1ms bin. The right plot displays the firing rate of individual neurons averaged across the simulation period.

cannot be distinguished from the background “distractor” (marked by blue) in terms of firing rate. A single post-synaptic neuron (membrane time constant $\tau_m=100\mu s$) integrates the input spikes via the memristive synapses and fires an action potential every time the membrane potential crosses over the defined threshold (membrane threshold $V_{th}=200$ in baseline case). The post-synaptic spikes are applied to the bottom electrode of all memristive synapses. The post-synaptic neuron is modelled as a leaky integrate and fire (LIF) neuron using Gerstner’s Spike Response Model [10]. The generation of the input spikes and the simulation of the post-synaptic neuron follow the description in [11].

The second-order memristor synapses naturally evolve in response to the applied pre- and post-synaptic spike trains. Simulation of the synaptic weight updates is based on the STDP function experimentally measured from devices with HI layers, shown in Fig. 1b. Specifically, a multiplicative STDP learning rule is used with the following parameters: $A_p=0.37$, $A_d=0.3$, $\tau_p=48.6\mu s$, $\tau_d=85.2\mu s$, $\eta=0.01$.

$$w = \begin{cases} w - \eta * (w - 0.2) * A_d * \exp\left(-\frac{\Delta t}{\tau_d}\right), & \Delta t < 0 \\ 0, & \Delta t = 0 \\ w + \eta * (1.0 - w) * A_p * \exp\left(-\frac{\Delta t}{\tau_p}\right), & \Delta t \geq 0 \end{cases} \quad (1)$$

B. Results

The initial weights of all synapses are set as 0.65 (normalized between 0-1), and the post-synaptic neuron fires periodically both in and outside the pattern duration in the initial stage, as shown in the top panel of Fig. 3b. This is due to the largely random spiking inputs to the post-neuron. The synaptic weights are then naturally updated according to the relative time between pre- and post-synaptic spikes following the observed STDP behavior. If the spike trains of the pre- and post-synaptic neurons do not have temporal correlation, the corresponding synaptic weight will gradually stabilize at the value where LTP and LTD balance out each other (~ 0.45 for

our device). If the pre-synaptic spikes contribute to the firing of the post-synaptic neuron, the corresponding synaptic weights will be strengthened. The next time the same pre-synaptic spiking pattern arrives, the enhanced synaptic




Fig. 3. (a) Post-synaptic spike latency, $t_{latency}$, as a function of post-neuron discharges, showing the selectivity emergence and latency reduction over time. (b) The membrane potential of the post-synaptic neuron in three stages: initial, middle and final. The grey boxes represent the pattern occurrence. The post-synaptic neuron transits from periodic firing with no selectivity in the initial stage to the selective firing only within the repeated pattern in the middle stage. From the middle to the final stage the post-synaptic neuron firing then tracks back to the beginning of the target pattern, further reducing $t_{latency}$.

weights will further promote the firing of the post-synaptic neuron, therefore forming a positive feedback loop to enhance the temporal causal relationship.

In the initial stage, the periodic post-synaptic spikes potentiate some afferent connections, and the selection is by chance. Among them, the connections whose corresponding input neurons fire the same input pattern repeatedly will then be gradually potentiated, causing the post-synaptic neuron to




Fig. 4. (a) Evolution of device conductance distribution throughout the simulation duration. Synapses whose input neurons are involved in the pattern are colored as magenta, and the other synapses whose input neurons are not involved in the pattern are colored as gray. (b) Raster plot of the 1,000 input neurons in the final 0.1ms of the simulation. The neurons are reordered so that neurons 0-499 are involved in the pattern and 500-999 are not involved. The plot is color coded based on the corresponding synaptic weights of the neurons. The grey rectangle marks the target spatiotemporal firing pattern. (c) The post-synaptic membrane potential induced by the input spikes shown in (b). The red vertical line and grey box mark the beginning and the duration of the spatiotemporal pattern, respectively.

fire more closely following the target pattern. Fig. 3 plots the post-synaptic spike latency with respect to the beginning of the target patterns ($t_{latency}$), as a function of the post-synaptic neuron discharges (firing). It is clearly shown that $t_{latency}$ changes from an initial random distribution to a concentrated band around 0.17ms. The selectivity emerges after around 1,000 post-synaptic firing events, when the temporally uncorrelated connections are depressed, and the correlated ones are strengthened. As shown in the middle panel of Fig.

3b, the post-synaptic neuron ceases to respond to the other “interfering” background spikes, and only respond to the target spatiotemporal pattern. Fig. 4a depicts the evolution of synaptic weights. The weights correlating to the pattern start to separate from the synaptic weights that are irrelevant to the pattern. As discussed earlier, the non-patterned weights settle down to ~ 0.45 after sufficient inputs. After the selectivity has been established, every time the post-synaptic neuron spikes within the pattern, the synapses of the afferents that fire slightly before the post-synaptic spike will be further potentiated, strengthening the causal relationship. Therefore, as the learning continues, the post-synaptic neuron tracks back through the pattern and t_{latency} decreases from around 0.17ms to 0.07ms, as shown in Figs 3a,b. The latency reduction will saturate when all the weights reach the maximum value.

The final stage of the simulation is shown in Fig. 4b and c. The input spikes in the last 0.1ms are color-coded according to the corresponding synaptic weights. The post-synaptic neuron membrane potential is plotted in Fig. 4c. The synaptic weights whose input neurons are involved in the early stage of the pattern have been potentiated, therefore contributing to a significant EPSP of the post-synaptic neuron and causing the post-synaptic neuron to fire, while those whose input neurons follow right after the post-synaptic spike are depressed to lower values. The spatiotemporal pattern is essentially a sequence of coincidence events. The post-synaptic neuron learns the temporal causality through the evolution of synaptic weights via STDP learning and acts as a coincidence detector. When the precise coincident events of input neurons arrive, the post-synaptic neuron will emit a spike, while it will remain quiet during equally dense background spikes.

C. Influence of Neural Network Parameters

Here, we investigate how the system performance can be affected by different network parameters. The evaluation is performed based on two factors: true positive rate (TPR) and number of false alarm spikes in the final 1.5s simulation duration. The first one is also referred to as the hit rate, which calculates the ratio of the patterns with post-synaptic spike and

the total number of patterns. The number of false alarm represents the number of post-synaptic spiking events outside the pattern.

As STDP is sensitive to time, the first parameter we check is the spike jittering. Fig. 5a shows that adding a 0.01ms jitter to the input spikes reduces the TPR from 91.8% to 83.1%, while false alarm remains zero up to 0.02ms jitter. Sharp deterioration of system performance is observed when the jitter increases from 0.03ms to 0.04ms. Considering the fact that the time constant of post-synaptic membrane potential is 0.1ms, a 0.04ms jitter can impair the temporal causality between the pre- and post-synaptic pulses, thus causing the failure of pattern recognition.

We also analyzed the effects of initial weights. In general, the system shows similar performance, where TPR is over 0.9 and the false alarm maintains near zero. The time it takes for the emergence of selectivity, however, will be affected by the initial weights. For the case of having initial weights ~ 0.5 , the post-synaptic neuron rapidly learns the pattern after only 250 discharges. This is due to the fact that the initial weights are closer to the stabilization value of the non-patterned weights (~ 0.45), therefore a few post-synaptic discharges can effectively suppress the non-patterned connections. While for higher initial weights as 0.8, it takes around 1,500 discharges before the initial selectivity is developed. In the end, the latency of the post-synaptic spike all converges to a similar value in the final stage of simulation, which means different initial weights have little influence on the system performance. We note that cases where all initial weights are set below 0.4 failed to converge, as the weights are too low to initiate post-synaptic events, therefore no weight updates can be induced.

Another parameter we studied is the portion of afferents involved in the pattern (p , 0.5 was used as the baseline case). With only 25% of afferents involved, the system can reliably detect the pattern with TPR over 80%. It is interesting to find that instead of increasing with p , TPR drops when p increased from 0.3 to 0.45. When $p = 0.3$, t_{latency} is 0.148ms, while t_{latency} is 0.092ms when the $p = 0.45$. Larger t_{latency} means the post-neuron has a longer time to integrate the weighted inputs. Additionally, longer t_{latency} means that a larger part of connections within the pattern will be potentiated, and fewer connections will be depressed. These effects combine together and lead to better performance at $p=0.35$ than $p=0.45$. It is not surprising to see that further increase of p beyond 0.45 leads to an improvement of the system performance, as more neurons repeatedly coincide to spikes within the pattern. The learning process will therefore be accelerated, and the selected weights can be potentiated to the maximum value within shorter time.

D. Device Non-idealities

Finally, the impact of different device non-idealities is evaluated. Standard deviations from 5% to 20% are added for each simulation parameter, which cover the range of experimentally measured variation in actual devices.

We first try to add variation to the initial weights, as the conductance update is dependent on the current device status. As shown in Fig. 6a, the initial weight variation has no impact on the system performance. As long as the simulation duration is long enough, the weights will gradually learn and reach the same distribution in the end.




Fig. 5. Impact of different neural network parameters. True positive rate (blue) and false alarm (orange) are evaluated as a function of (a) Jitter, (b) Initial weight, (c) Percentage of spike deleted within the pattern, and (d) Proportion of afferents involved in the pattern.

The amplitude (A_p/A_d) and time constant (τ_p/τ_d) variations among devices in STDP updates are investigated in Figs. 6b and c. The system shows no obvious deterioration in performance up to 20% variation. This can be explained by the fact that the pattern involves the firing events of a population of neurons instead of a single neuron. The device-to-device variations are essentially averaged out by the population. As the post-synaptic neuron integrates the inputs from all afferents, it is not vulnerable to the variations in individual devices.




Fig. 6. Impact of device non-idealities. True positive rate (blue) and false alarm (orange) are evaluated considering the influence of (a) Initial weight variation, (b) Device-to-device amplitude variation, (c) Device-to-device time constant variation, (d) Cycle-to-cycle variation.

As the conductance update of memristive devices is stochastic in nature, we also evaluate cycle-to-cycle variations. The system performance, as plotted in Fig. 6d, is not affected. The weight change of devices for each spiking event is essentially small, which means the learning is progressive and cumulative. The final weights are determined by the accumulation of thousands of discharges, where cycle-to-cycle variation gets cancelled out.

These evaluations confirm the robustness of the proposed system to the inherent device non-idealities. The memristor-based system can thus reliably act as a coincidence detector even with considerable amount of variations.

IV. CONCLUSION

Memristors are promising candidates to construct compact and energy-efficient neuromorphic systems. By taking advantage of the intrinsic dynamics of second-order state variables, 2nd-order memristors can natively implement STDP learning rule in a bio-plausible manner. We demonstrate that a single post-synaptic neuron when combined with second-order memristive synapses can learn to detect a spatio-temporal spiking pattern covered by interfering background spikes at similar firing rates. The STDP rule selectively potentiates the synapses of the input neurons whose firing

events contribute to the post-synaptic neuron firing and depresses the synapses of the input neurons that tend to fire right after the post-synaptic spike. The synapses of the other input neurons whose spike trains has no temporally correlation with the post-synaptic spikes gradually evolve to a value where LTP and LTD are balanced. As the learning proceeds, the post-synaptic neuron develops selectivity and tracks back to the beginning of the pattern. The system also shows high robustness to device non-idealities and different neural network parameters.

ACKNOWLEDGEMENTS

The authors would like to thank Dr. S. H. Lee and Dr. M. Zidan for helpful and stimulating discussions. This work was supported in part by the National Science Foundation through awards ECCS-1915550 and DMR-1810119.

REFERENCES

- [1] Kim, S., Du, C., Sheridan, P., Ma, W., Choi, S., and Lu, W.D., "Experimental demonstration of a second-order memristor and its ability to biorealistically implement synaptic plasticity", *Nano Lett.*, 2015, 15, (3), pp. 2203-2211
- [2] Tuma, T., Le Gallo, M., Sebastian, A., and Eleftheriou, E., "Detecting correlations using phase-change neurons and synapses", *IEEE Electron Device Lett.*, 2016, 37, (9), pp. 1238-1241
- [3] Sebastian, A., Tuma, T., Papandreou, N., Le Gallo, M., Kull, L., Parnell, T., and Eleftheriou, E., "Temporal correlation detection using computational phase-change memory", *Nat. Commun.*, 2017, 8, (1)
- [4] Wu, Y., Moon, J., Zhu, X., and Lu, W.D., "Neural Functional Connectivity Reconstruction with Second - Order Memristor Network", *Adv. Intell. Syst. Comput.*, pp. 2000276.
- [5] Moon, J., Wu, Y., Zhu, X., and Lu, W.D., "Neural connectivity inference with spike-timing dependent plasticity network", *Science China Information Sciences*, 2021, 64, (6), pp. 1-10
- [6] Prezioso, M., Mahmoodi, M., Bayat, F.M., Nili, H., Kim, H., Vincent, A., and Strukov, D., "Spike-timing-dependent plasticity learning of coincidence detection with passively integrated memristive circuits", *Nat. Commun.*, 2018, 9, (1), pp. 1-8
- [7] Nádasdy, Z., Hirase, H., Czúrkó, A., Csicsvari, J., and Buzsáki, G., "Replay and time compression of recurring spike sequences in the hippocampus", *Journal of Neuroscience*, 1999, 19, (21), pp. 9497-9507
- [8] Reinagel, P., and Reid, R.C., "Precise firing events are conserved across neurons", *Journal of Neuroscience*, 2002, 22, (16), pp. 6837-6841
- [9] Fellous, J.-M., Tiesinga, P.H., Thomas, P.J., and Sejnowski, T.J., "Discovering spike patterns in neuronal responses", *Journal of Neuroscience*, 2004, 24, (12), pp. 2989-3001
- [10] Masquelier, T., Guyonneau, R., and Thorpe, S.J., "Spike timing dependent plasticity finds the start of repeating patterns in continuous spike trains", *PloS one*, 2008, 3, (1), pp. e1377
- [11] Gerstner, Wulfram, and Werner M. Kistler, "Spiking neuron models: Single neurons, populations, plasticity", Cambridge university press, 2002.
- [12] Jo, S.H., Chang, T., Ebong, I., Bhadviya, B.B., Mazumder, P., and Lu, W., "Nanoscale memristor device as synapse in neuromorphic systems", *Nano Lett.*, 2010, 10, (4), pp. 1297-1301
- [13] Yu, S., Wu, Y., Jeyasingh, R., Kuzum, D., and Wong, H.-S.P., "An electronic synapse device based on metal oxide resistive switching memory for neuromorphic computation", *IEEE Trans. Electron Devices*, 2011, 58, (8), pp. 2729-2737
- [14] Alibart, F., Pleutin, S., Bichler, O., Gamrat, C., Serrano - Gotarredona, T., Linares - Barranco, B., and Vuillaume, D., "A memristive nanoparticle/organic hybrid synapstor for neuroinspired computing", *Adv. Funct. Mater.*, 2012, 22, (3), pp. 609-616
- [15] Serrano-Gotarredona, T., Masquelier, T., Prodromakis, T., Indiveri, G., and Linares-Barranco, B., "STDP variations with memristors for spiking neuromorphic learning systems", *Frontiers in neuroscience*, 2013, 7, pp. 2