

# Memristive Reservoirs Learn to Learn

Ruomin Zhu  
School of Physics,  
The University of Sydney  
Sydney, NSW, Australia  
rzhu0837@uni.sydney.edu.au

Jason K. Eshraghian  
Department of Electrical and  
Computer Engineering,  
University of California, Santa Cruz  
Santa Cruz, CA, USA  
jeshragh@ucsc.edu

Zdenka Kuncic  
School of Physics and  
Sydney Nano Institute,  
The University of Sydney  
Sydney, NSW, Australia  
zdenka.kuncic@sydney.edu.au

## ABSTRACT

Memristive reservoirs draw inspiration from a novel class of neuromorphic hardware known as nanowire networks. These systems display emergent brain-like dynamics, with optimal performance demonstrated at dynamical phase transitions. In these networks, a limited number of electrodes are available to modulate system dynamics, in contrast to the global controllability offered by neuromorphic hardware through random access memories. We demonstrate that the learn-to-learn framework can effectively address this challenge in the context of optimization. Using the framework, we successfully identify the optimal hyperparameters for the reservoir. This finding aligns with previous research, which suggests that the optimal performance of a memristive reservoir occurs at the ‘edge of formation’ of a conductive pathway. Furthermore, our results show that these systems can mimic membrane potential behavior observed in spiking neurons, and may serve as an interface between spike-based and continuous processes.

## CCS CONCEPTS

• **Computing methodologies** → **Supervised learning by regression; Multi-task learning;** • **Theory of computation** → **Optimization with randomized search heuristics.**

## KEYWORDS

neuromorphic, memristive, reservoir, learn-to-learn, meta-learning, spiking neurons

### ACM Reference Format:

Ruomin Zhu, Jason K. Eshraghian, and Zdenka Kuncic. 2023. Memristive Reservoirs Learn to Learn. In *Proceedings of International Conference on Neuromorphic Systems 2023 (ICONS 2023)*. ACM, New York, NY, USA, 7 pages. <https://doi.org/XXXXXXXX.XXXXXXX>

## 1 INTRODUCTION

Nanowire networks are a novel class of neuromorphic devices that demonstrate potential for brain-inspired computing and information processing [5, 15, 21–25, 30, 35, 36, 40–42]. By embedding memristive switching dynamics into naturally-arising neuromorphic connectivity structures [26, 29], these networks display emergent

brain-like dynamics [6], including dynamical phase transitions and avalanche criticality [8, 15].

The synaptic sites of nanowire networks are not directly accessible, in contrast to random access memories (RAM) [4, 9, 17], where each memory cell is addressable and programmable. The lack of controllability is compensated for by the dynamic nature of nanowire networks, which is a key feature that enables them to adapt to evolving input signals. Nevertheless, it is worth investigating how these neuromorphic systems can be optimized for information processing tasks. For example, previous studies have shown that in a physical reservoir computing framework, nanowire networks can achieve superior learning performance when operating near a dynamical phase transition [15, 40]. Rather than manually exploring the optimal region of operation, an alternative and more effective way to optimize the parameters that can be physically adjusted (namely, inputs and outputs) is the learn-to-learn (L2L) framework, commonly known as meta-learning.

The L2L approach is a scheme for optimizing learning capacity from prior experiences [16]. Recent advances of L2L algorithms are built upon the premise that the learning system – typically an artificial neural network – is differentiable, so that the gradients can be propagated through the network to adjust the hyperparameters (e.g. learning rate) [1, 11, 14]. It has also been shown that biologically or physically inspired systems are suitable candidates as learning agents in the L2L framework [2, 37]. In particular, Bohnstingl et al. [3] showed that non-differentiable systems could be optimized using non-gradient-based optimization schemes.

This study demonstrates how learning is achieved using the collective dynamics of a system abstracted from physical nanowire networks (see details in [15, 40]) under a physical reservoir computing (RC) framework [18, 28], namely a memristive reservoir, where training is restricted to the readout layer to circumvent the computation burden of training traditional deep artificial network architectures [27]. The synaptic sites (recurrent weights) in the reservoir are not individually programmable but this may be offset by the highly rich set of dynamics available in the memristive substrate as a response to external stimuli. The L2L approach is applied to the memristive reservoir and we show that this system is able to learn the dynamics of a family of nonlinearly filtered signals. Furthermore, we also demonstrate that the memristive reservoir system is able to generate dynamics that resemble the membrane potentials of spiking neural networks (SNNs) directly from continuous inputs, which implies its potential to bridge the gap between continuous signals and spike-based computing paradigms. By learning across various membrane potential dynamics, meta-learning can identify common principles or shared features that govern the dynamics of both spiking neurons and nanowire networks, which

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

ICONS 2023, Aug 1–3, 2023, Santa Fe, NM, USA

© 2023 Association for Computing Machinery.

ACM ISBN 978-1-4503-XXXX-X/18/06...\$15.00

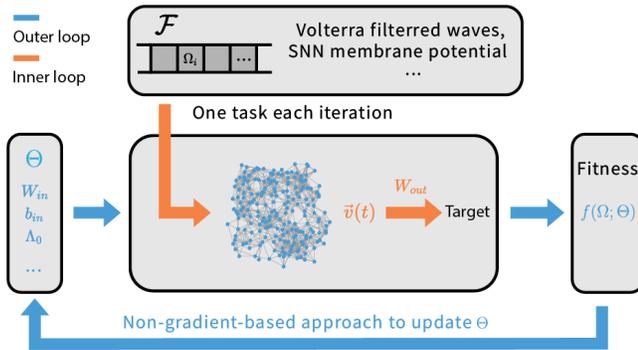
<https://doi.org/XXXXXXXX.XXXXXXX>

may offer deeper insights into the underlying ionic mechanisms that pervade both biology and emerging memory technologies. In particular, a nested-loop structure is implemented as following:

- L2L is used in the outer loop to determine the optimal hyperparameters of the memristive reservoir;
- Linear regression is employed in the inner loop to optimize the stimulus/readout protocol to achieve the task objectives.

Demonstrating how L2L is compatible with systems grounded in physics, such as memristive reservoirs, opens up the potential to optimize learning using adaptive networks where individual internal weights are not trained, but rather allowed to self-adjust in response to dynamical inputs, in a manner similar to the brain’s neural network. This has implications beyond memristive reservoirs and could lead to more efficient and effective optimization schemes for learning complex physical phenomena.

## 2 RESEARCH METHODS



**Figure 1: Schematic diagram of the L2L framework. At each iteration of the outer loop, the hyperparameter set  $\Theta = \{W_{in}, b_{in}, \Lambda_0, \dots\}$  is perturbed in multiple directions. These perturbed hyperparameters are passed to the inner loop and learning task  $\Omega$  randomly selected from  $\mathcal{F}$  are performed.  $\Theta$  is then optimized in the outer loop based on the returned fitness of different tasks. The inner loop operations are indicated by orange arrows and the outer loop operations by blue arrows.**

As continuous temporal input voltage signals are delivered to a memristive reservoir via dedicated input nodes, the network autonomously adjusts its internal state. Voltages of a subset of the remaining nodes can be read out and used to train weights in the fully-connected linear output layer. Two L2L task families are studied: learning nonlinear Volterra dynamics and learning membrane dynamics from SNNs.

### 2.1 The L2L Framework

Fig. 1 illustrates the L2L framework, where two iterative loops work together to optimize the system’s learning performance for a chosen family of tasks  $\mathcal{F}$ . In this work, the task families are split into two subsets, one used for meta-training and the other for meta-testing.

The objective of the outer loop (orange in Fig. 1) is to optimize the reservoir hyperparameters  $\Theta$  based on the fitness in the inner loop.

In each iteration of the inner loop (blue in Fig. 1), the system learns from a specific task  $\Omega_i$  within the meta-training phase of  $\mathcal{F}$  and the fitness  $f(\Omega_i; \Theta)$  is evaluated. Non-gradient-based techniques, namely simulated annealing (SA) and evolutionary strategies (ES), are applied to the fitness to determine the optimal hyperparameter set  $\Theta'$  (or range), thus achieving the best collective learning outcome for  $\mathcal{F}$ :

$$\Theta' = \underset{\Theta}{\operatorname{argmax}}(f(\Omega; \Theta)), \quad \Omega \in \mathcal{F}. \quad (1)$$

After 100 generations of outer-loop training, the performance of the baseline (without meta-learning) and meta-learned reservoirs are evaluated using tasks from the meta-testing set to evaluate the influence of the L2L algorithm on the reservoir’s learning performance.

**2.1.1 Simulated Annealing.** Simulated annealing optimizes the objective function by mimicking the physical annealing process [20]. The algorithm is parameterized by a decaying ‘temperature’  $T$ . At each generation of the outer loop, the hyperparameters are perturbed in multiple directions by a step size  $\epsilon$  drawn from a normal distribution parameterized by  $T$ . For each perturbed hyperparameter  $\tilde{\Theta}$ , the change in fitness ( $\Delta f = f(\Omega, \tilde{\Theta}) - f(\Omega; \Theta)$ ) is estimated.  $\tilde{\Theta}$  is accepted if the corresponding fitness improves, otherwise the probability ( $P$ ) of accepting  $\tilde{\Theta}$  is determined by an exponential distribution:

$$P = \begin{cases} 1, & \Delta f \geq 0, \\ e^{\Delta f/T}, & \Delta f < 0. \end{cases} \quad (2)$$

When  $T$  is high, some sub-optimal solutions can be accepted and the algorithm explores a broader parameter space. As it ‘cools down’, the probability of accepting worse solutions decreases and the solution eventually converges if a global optimum exists.

**2.1.2 Evolution Strategies.** Evolution strategies are inspired by natural selection and evolution [31, 38]. A population of candidate solutions are generated and evolved based on weighted fitness and a pre-determined learning rate  $\eta$  [32, 34]. At generation  $k$ , each candidate  $\Theta_{k,i}$  in the population is perturbed by a step size  $\epsilon_i$  drawn from a normal distribution  $(0, \mathbb{I})$ , with standard deviation  $\sigma$

After evaluating the fitness  $f(\Omega_i; \Theta_{k,i} + \sigma\epsilon_i)$  for all candidates in the same generation,  $\Theta_{k,i}$  is evolved by:

$$\Theta_{k+1,i} = \Theta_{k,i} + \frac{\eta}{n\sigma} \sum_{i=1}^n f(\Omega_i; \Theta_{k,i} + \sigma\epsilon_i)\epsilon_i, \quad (3)$$

where  $n$  is the size of the population. The statistical setup enables ES to deal with high-dimensional problems, making it useful for optimization problems with many parameters.

### 2.2 Learning Volterra dynamics

Consider a two-terminal configuration of a memristive reservoir with one source node and one drain node (see [40] for details). Fig. 2 shows how the memristive reservoir is used to learn the nonlinear Volterra dynamics (Task 1). As the green box and arrows indicate, the input signal  $x(t)$  is delivered to the source node, while the drain node is grounded and an external fully connected output layer linearly combines the voltage read outs  $\vec{v}(t)$  of 64 other nodes to regress to each target signal  $u(t)$ . The objective of the L2L process

here is to fine-tune the hyperparameters to achieve optimal regression for the nonlinear time-delayed target signal. The input weight ( $W_{in}$ ), input bias ( $b_{in}$ ), and initial reservoir state ( $\Lambda_0$ ) are optimized by L2L via simulated annealing.  $\Lambda_0$  is parameterized by a pulse of 1 V DC with varying width  $T_0$  applied to the input node prior to the task, as described in [40].

The input signal  $x(t)$  is generated by combining two sine waves:

$$x(t) = A_1 \sin(2\pi \frac{t}{T_1} + \phi_1) + A_2 \sin(2\pi \frac{t}{T_2} + \phi_2), \quad (4)$$

where  $T_1 = 0.323$  s and  $T_2 = 0.5$  s, and where  $A_1, A_2 \in [0.5, 1]$  and  $\phi_1, \phi_2 \in [0, \frac{\pi}{2}]$  are chosen randomly. The target signal  $u(t)$  is generated using a second-order Volterra filter:

$$u(t) = \int_{\tau} k_{\Omega}^1(\tau)x(t-\tau)d\tau + \int_{\tau_1} \int_{\tau_2} k_{\Omega}^2(\tau_1, \tau_2)x(t-\tau_1)x(t-\tau_2)d\tau_1d\tau_2, \quad (5)$$

in which  $\tau, \tau_1, \tau_2 \in [0, 0.5]$  denote the time delay of the signals, while  $k_{\Omega}^1$  and  $k_{\Omega}^2$  are task-specific random Volterra kernels (see [37] for further information on generating the kernels). 100 target signals are generated using different Volterra kernels to comprise the family  $\mathcal{F}$ . The green box in Fig. 2 shows an example of one input signal (black) and Volterra-filtered target signals (colored) with different random kernels.

The readout weights  $W_{out}$  are trained using least square to minimize the loss:

$$\mathcal{L} = \|U - W^T V\|_2, \quad (6)$$

where  $\|\cdot\|_2$  represents the L-2 norm, and  $U = [u(1), u(2), \dots, u(t)]$  and  $V = [\vec{v}(1), \vec{v}(2), \dots, \vec{v}(t)]$  are the stacked target and readout of the reservoir. The learning results are estimated by normalized root mean squared error (NRMSE):

$$\text{NRMSE} = \frac{\sqrt{\frac{1}{T} \sum_{t=1}^T (W_{out}^T \vec{v}(t) - u(t))^2}}{u_{max} - u_{min}}, \quad (7)$$

where  $T$  is the length of the signal.

### 2.3 Learning membrane dynamics of SNNs

The blue arrows in Fig. 2 illustrate the scheme for this task family. A set of 20 randomly generated Volterra filter signals  $\vec{u}(t)$  are utilized as input and converted to spike trains  $\vec{u}'(t)$  using delta modulation. A fully connected SNN ( $\mathcal{S}$ , developed using [10]) with 20 input neurons, 1,000 hidden neurons, and 5 output neurons receives  $\vec{u}'(t)$  as input. For each task, the internal weights  $W_{\mathcal{S}}$  of the SNN are randomly generated, and the membrane potentials of the 5 output neurons,  $\vec{m}(t)$  are employed as the target signals of the learning task:

$$\vec{m}(t) = \mathcal{S}(\vec{u}'(t); W_{\mathcal{S}}). \quad (8)$$

The continuous Volterra signals  $\vec{u}(t)$  are delivered to 20 input nodes in the memristive reservoir ( $\mathcal{R}$ ), and 64 nodes are used as voltage readouts:

$$\vec{v}(t) = \mathcal{R}(W_{in} \odot \vec{u}(t) + \vec{b}_{in}; \Lambda_0). \quad (9)$$

Notice that  $W_{in}$  and  $\vec{u}(t)$  have the same dimensions and  $\odot$  represents element-wise multiplication.

For each task, the entire data stream is divided into three parts: the first 2000 data points are considered as a transient phase [19]; the subsequent 6000 data points (from 2000th to 8000th) employed as the support set (inner-loop training), and the last 1000 data points (from 8000th to 9000th) comprise the query set (inner-loop testing). The readout weights  $W_{out}$  are trained for each task using ridge regression to minimize the loss function:

$$\mathcal{L} = \|M - W_{out}^T V\|_2^2 + \alpha \|W_{out}\|_2^2, \quad (10)$$

where  $\alpha$  is a regularization term ranging between  $10^{-4}$  and 1, with  $M = [\vec{m}(1), \vec{m}(2), \dots, \vec{m}(t)]$  representing the stacked target membrane potentials. A 5-fold cross validation scheme is employed to avoid overfitting [13]. NRMSE of the query regime is reported as result.

A family  $\mathcal{F}$  of 150 tasks is created by utilizing SNNs with different internal weights. The goal of L2L in this context is to adapt the reservoir to new tasks by extrapolating information from past learning experiences, allowing it to generate desired membrane potentials. Simulated annealing and evolution strategies are used separately to optimize the input weight ( $W_{in}$ ) and the input bias ( $\vec{b}_{in}$ ) of each signal.  $\Lambda_0$  is fixed to the phase transition regime found in task 1.

## 3 RESULTS

### 3.1 L2L Volterra dynamics

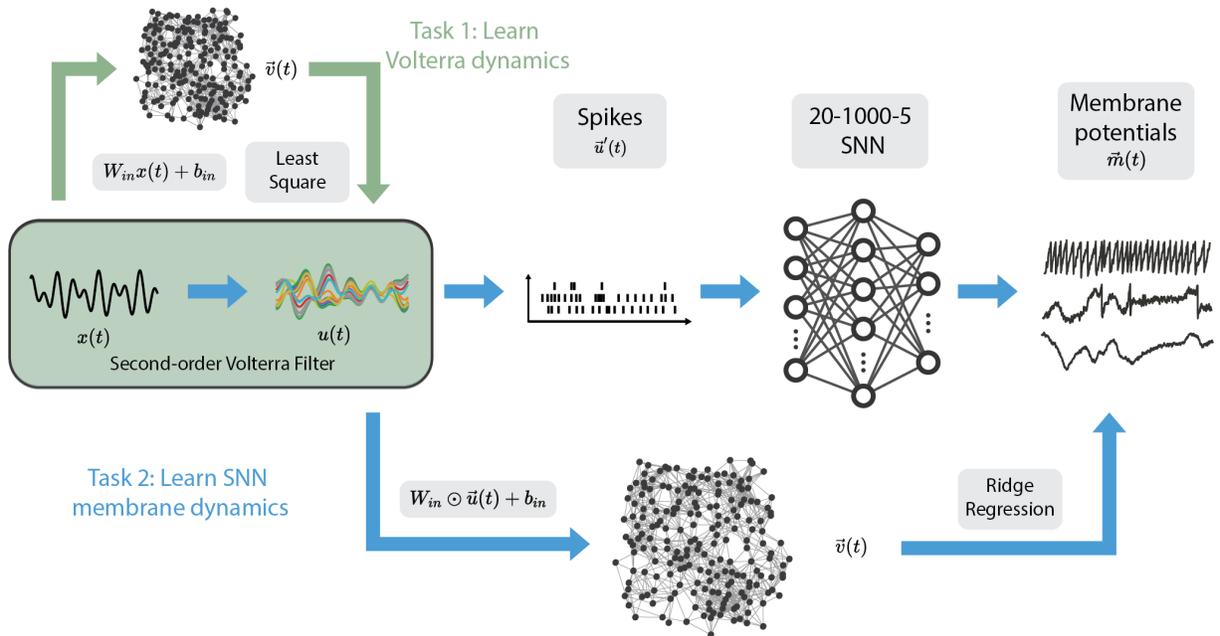
**Table 1: NRMSE for Task 1 with and without meta-learning**

	Meta-train	Meta-test
w/o meta-learning	0.209±0.003	0.211±0.014
Meta-learned (SA)	0.168±0.008	0.164±0.009

Fig. 3(a) shows the reservoir’s conductance (blue curve) in response to a 1 V DC pulse of varying width  $T_0$ . The shaded region represents the general phase transition regime, identified from previous studies [40], and the dashed line at  $T_0 = 2.3$  s indicates the optimal initial reservoir state  $\Lambda_0$  found by the SA optimization scheme, as shown in Fig. 3(b).  $T_0$  (orange) converges toward  $\approx 2.3$  s as RNRMSE (green) converges to a minimum after approximately 60 generations. Fig. 3(c) compares the system’s learning outcomes for a specific task in the Volterra family before (blue) and after (red) the outer loop optimization (cf. Fig. 1) and Table. 1 shows the corresponding NRMSE of the whole meta-testing set. A notable improvement in regression to the target signal is evident.

These results demonstrate that the L2L framework is successful in fine-tuning the reservoir’s initial state  $\Lambda_0$ . Remarkably, the optimization scheme finds the optimal  $\Lambda_0$  coinciding with the dynamical phase transition region identified independently in previous studies [40] as producing optimal task performance from memristive reservoirs.

To gain deeper insight into the internal dynamics of memristive reservoirs, Fig. 4 shows visualisation snapshots of network graphs of the reservoir for different  $T_0$ . When the reservoir is under-activated ( $T_0 < 2$  s), most memristive components are inactive and not enough information can be extracted to perform learning. On the other hand, when the reservoir is over-activated ( $T_0 > 8$  s), the



**Figure 2: Schematic diagram for the learning tasks employed in this study. The green box and arrows summarize the learning Volterra dynamics task family (Task 1), where the original signal  $x(t)$  is the input and the Volterra-filtered signal  $u(t)$  is the target for the reservoir. The blue arrows illustrate the flow of the learning SNN dynamics task family (Task 2), in which a set of 20 Volterra-filtered signals  $\bar{u}(t)$  is converted to spikes  $\bar{u}'(t)$  and delivered to a fully connected SNN. The membrane potentials  $\bar{m}(t)$  of the readout neurons are used as target while the set of continuous Volterra-filtered signals are delivered directly to the reservoir as input.**

internal dynamics saturates. The initial reservoir state at  $T_0 \approx 2.17$  s results in the best task performance and qualitatively, it is evident from Fig. 4 that this corresponds to an intermediate state, where conductance paths first span the network. At this ‘edge of formation’, the internal state of the reservoir produces voltage readout features that are more diverse than at later activation times, as shown in the corresponding node voltage distributions in Fig. 5.

### 3.2 L2L SNN dynamics

**Table 2: NRMSE for Task 2 with and without meta-learning**

	Meta-train	Meta-test
w/o meta-learning	0.212±0.041	0.235±0.036
Meta-learned (SA)	0.139±0.011	0.164±0.012
<b>Meta-learned (ES)</b>	<b>0.109±0.011</b>	<b>0.128±0.004</b>

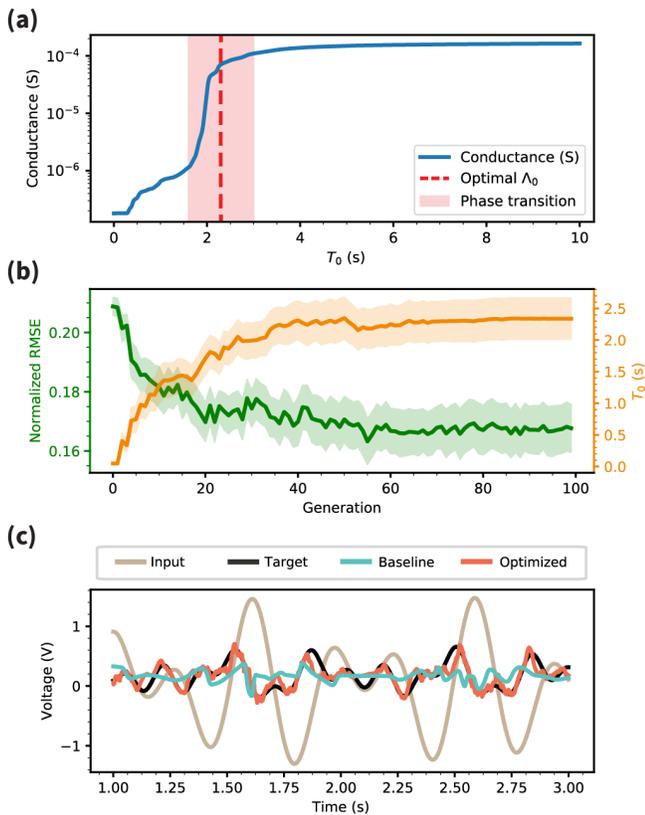
Fig. 6(a) shows the NRMSE of the query part (task-specific inner loop testing) for Task 2 during outer loop training for the two gradient-free optimization strategies considered. Similar to the previous task family, it can be observed that learning outcome improves with number of training generations. Fig. 6(b) compares the readouts generated by a baseline reservoir and an ES-optimized reservoir to the target curves (SNN membrane potentials) in a single meta-testing task. As is evident from the query period, the optimized system is considerably better in reproducing the fluctuating

dynamics of the membrane potential, which governs the spiking behaviors of the SNN readout neurons.

To assess the learning gain from the L2L process, Table 2 compares the NRMSE for Task 2 with and without meta-learning and shows substantial improvement with optimized reservoirs, up to  $\approx 50\%$  in the case of ES optimization. Table 2 and Fig. 6(a) also indicate that the learning outcome converges to a better result with ES compared to SA optimization. This is because this task family (Task 2) employs more hyperparameters than the Volterra task family (Task 1) and ES typically performs better for a larger parameter space while SA is more suitable for stepping out of local minima.

## 4 DISCUSSION

Previous studies have demonstrated the memory capacity of memristive reservoirs as well as their capability to generate dynamical features [12, 35, 40]. The results here suggest that memristive reservoirs are able to combine these properties together to enable learning of the rich, nonlinear time-delayed dynamics embedded by the Volterra filter. Furthermore, the L2L framework consistently determined that the optimal initial state for the memristive reservoir was at what can be described as the ‘edge of formation’. This point exists between where the memristive components first become activated (cf. Fig. 4,  $T_0=2.0$  s), and before an exponential cascade of parallel paths form (cf. Fig. 4,  $T_0=2.3$  s). The conductance exponentially ramps up at the time of initial formation, and saturates as more parallel pathways form. The intermediate internal state,

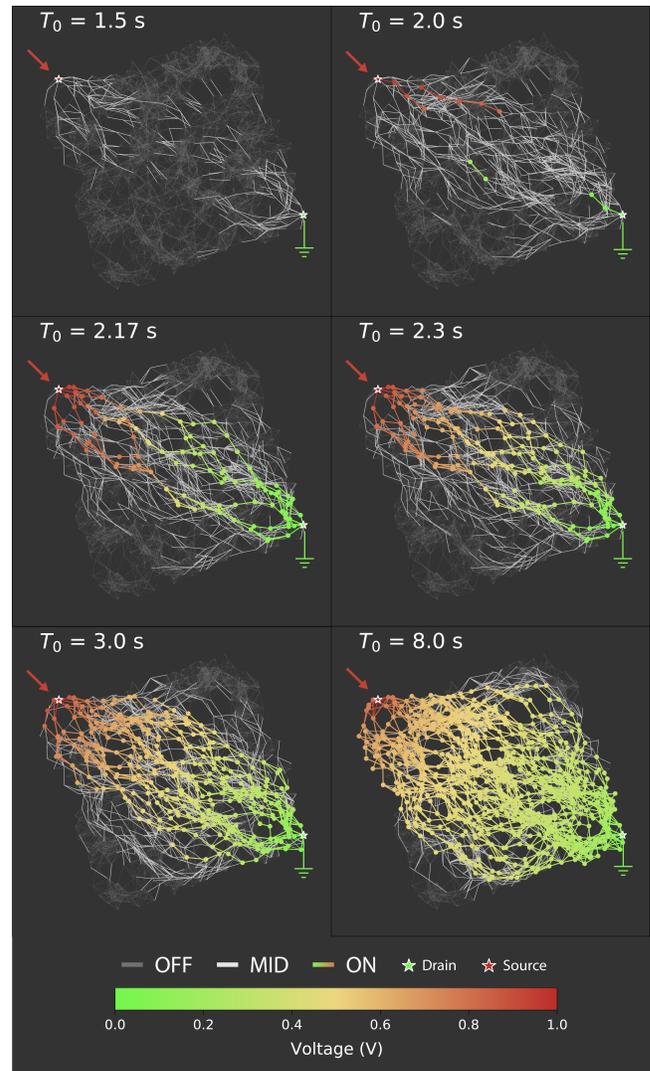


**Figure 3: L2L Volterra dynamics.** (a) Memristive reservoir conductance response (blue) to a varying 1 V DC pulse width  $T_0$ , showing a characteristic phase transition regime (shaded). The dashed red line represents the reservoir’s optimal initial state for learning Volterra dynamics, as found by the L2L scheme. (b) Normalized RMSE of the learning tasks and optimal  $T_0$  with respect to outer loop generation number. (c) Input signal, corresponding target, and the resulting learning outcomes for one representative task in the family (baseline is before optimization).

i.e., the ‘edge of formation’ was meta-learned as the optimal starting point prior to training in the outer loop.

This result opens up deeper insights to how memristive reservoirs can be optimally used in computation. It is somewhat intuitive that the reservoir does not perform at the time of initial formation, because the exponentially ramping conductance increase is highly unstable and challenging to control. The opposite problem exists after conductive pathways are formed, where switching has a negligible impact on the dynamics of the nanowire network. The ‘edge of formation’ can be thought of as a linear region in small-signal analysis, that fosters a controllable learning environment optimal for learning the rich dynamics of higher-order systems.

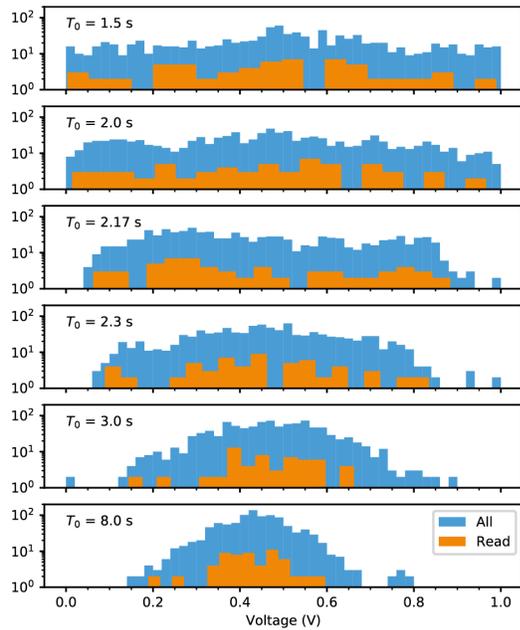
Additionally, this work also demonstrates that memristive reservoirs can learn the fluctuating dynamics of SNN membrane potentials. This is possible because, as shown in a previous study [15],



**Figure 4: Reservoir activity with different pre-initialization pulse widths.** A 1 V DC signal of varying width is delivered to the reservoir in prior to the task. Nodes on the conductance path from source to drain are colored based on their voltages. Edges are categorized as high (ON), intermediate (MID) and low (OFF) conductance levels and colored respectively.

memristive switching in a heterogeneous, recurrent network produces fluctuating dynamics that resemble action potentials which, when thresholded, generate spikes. In other words, the internal dynamics of memristive reservoirs effectively encode spike-like features into continuous input signals.

An immediate implication of this result is that memristive reservoirs have the potential to serve as an interface between continuous data streams and spike-based computing paradigms, which could substantially improve the workflow for SNN applications. The use of trainable spike-based embeddings has offered an alternative approach to classical rate and temporal encoders in many recent SNN



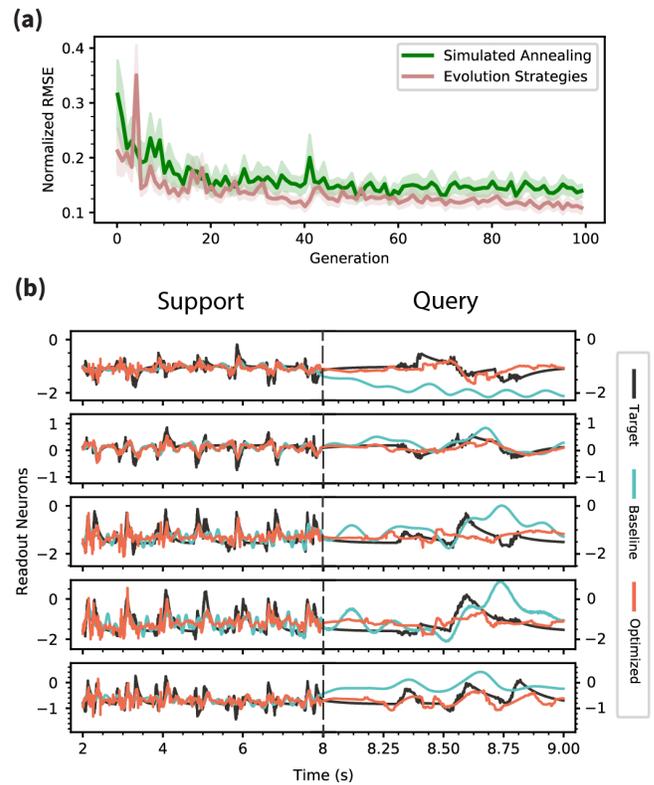
**Figure 5: Reservoir node voltage histograms for different pre-initialization pulse widths  $T_0$ . Blue indicates all nodes, orange indicates nodes used as readouts for the learning tasks.**

works, but relies on gradient-based optimization to compress data into more efficient spike-based representations [7, 39, 43]. These results demonstrate how native internal dynamics, constrained by physics, naturally give rise to embeddings that can ultimately reconstruct signals in the context of gradient-free meta-learning.

## 5 CONCLUSION

This study shows that the L2L framework can effectively adjust the hyperparameters of memristive reservoirs, attaining a similar optimal regime as found by a manual search in previous studies. Moreover, we demonstrated that the learning capability of the system can be extended and optimized for a family of related tasks, rather than being limited to a single task. This approach could pave the way for highly adaptive learning tasks based on real-world settings, similar to how biological brains can learn quickly from limited examples. Our finding that memristive reservoirs have the ability to reproduce SNN membrane potentials is significant because of the potential to use them in place of spike-based embeddings or encoders. In particular, membrane potential dynamics in complex tasks may be encoded using memristive reservoirs as a more efficient mode of compression.

The memristive reservoirs studied here are motivated by physical reservoirs using self-assembled nanowire networks. The neuromorphic properties of these networks are influenced by various factors such as nanowire density, diameter, average length and amount of dielectric. Similar to how evolutionary optimization schemes can be utilized to design and train neuromorphic systems [33], it may be possible in the future to exploit the L2L framework to find the



**Figure 6: L2L SNN membrane dynamics. (a) Normalized RMSE of the learning tasks as a function of generations in the outer loop. (b) Membrane potentials of 5 readout neurons from the SNN (black) overlaid by the learned dynamics from the baseline (blue) and the optimized (orange) memristive reservoirs for one representative meta-testing task in the task family. The left and right columns of panels show results for the support and query periods of the task, respectively. Note: a different timescale is used during the query phase for better visibility at test time.**

optimal nanowire networks for specific task families, effectively realizing a ‘learn-to-build’ approach.

## ACKNOWLEDGMENTS

The authors acknowledge the complimentary computing resources provided by Google Cloud. The authors would like to thank Anand Subramoney for inspirational discussions on the L2L framework. R.Z. is supported by the PREA scholarship from the University of Sydney.

## REFERENCES

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gomez, Matthew W. Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. 2016. Learning to Learn by Gradient Descent by Gradient Descent. <https://doi.org/10.48550/arXiv.1606.04474> arXiv:arXiv:1606.04474
- [2] Guillaume Bellec, Darjan Salaj, Anand Subramoney, Robert Legenstein, and Wolfgang Maass. 2018. Long Short-Term Memory and Learning-to-Learn in Networks of Spiking Neurons. arXiv:arXiv:1803.09574

- [3] Thomas Bohnstingl, Franz Scherr, Christian Pehle, Karlheinz Meier, and Wolfgang Maass. 2019. Neuromorphic Hardware Learns to Learn. *Frontiers in Neuroscience* 13 (2019).
- [4] Ting-Chang Chang, Kuan-Chang Chang, Tsung-Ming Tsai, Tian-Jian Chu, and Simon M. Sze. 2016. Resistance Random Access Memory. *Materials Today* 19, 5 (June 2016), 254–264. <https://doi.org/10.1016/j.mattod.2015.11.009>
- [5] E. C. Demis, R. Aguilera, H. O. Sillin, K. Scharnhorst, E. J. Sandouk, M. Aono, A. Z. Stieg, and J. K. Gimzewski. 2015. Atomic Switch Networks—Nanoarchitectonic Design of a Complex System for Natural Computing. *Nanotechnology* 26, 20 (April 2015), 204003. <https://doi.org/10.1088/0957-4484/26/20/204003>
- [6] Adrian Diaz-Alvarez, Rintaro Higuchi, Paula Sanz-Leon, Ido Marcus, Yoshitaka Shingaya, Adam Z. Stieg, James K. Gimzewski, Zdenka Kuncic, and Tomonobu Nakayama. 2019. Emergent Dynamics of Neuromorphic Nanowire Networks. *Scientific Reports* 9, 1 (Dec. 2019), 14920. <https://doi.org/10.1038/s41598-019-51330-6>
- [7] Dominik Dold. 2022. Relational representation learning with spike trains. In *2022 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 1–8.
- [8] Christopher S Dunham, Sam Lilak, Joel Hochstetter, Alon Loeffler, Ruomin Zhu, Charles Chase, Adam Z Stieg, Zdenka Kuncic, and James K Gimzewski. 2021. Nanoscale Neuromorphic Networks and Criticality: A Perspective. *Journal of Physics: Complexity* 2, 4 (Dec. 2021), 042001. <https://doi.org/10.1088/2632-072X/ac3ad3>
- [9] Jason K Eshraghian, Xinxin Wang, and Wei D Lu. 2022. Memristor-based binarized spiking neural networks: Challenges and applications. *IEEE Nanotechnology Magazine* 16, 2 (2022), 14–23.
- [10] Jason K. Eshraghian, Max Ward, Emre Neftci, Xinxin Wang, Gregor Lenz, Girish Dwivedi, Mohammed Benmamoun, Doo Seok Jeong, and Wei D. Lu. 2021. Training Spiking Neural Networks Using Lessons From Deep Learning. (Sept. 2021). <https://doi.org/10.48550/arXiv.2109.12894>
- [11] Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. arXiv:arXiv:1703.03400
- [12] K. Fu, R. Zhu, A. Loeffler, J. Hochstetter, A. Diaz-Alvarez, A. Stieg, J. Gimzewski, T. Nakayama, and Z. Kuncic. 2020. Reservoir Computing with Neuromemristive Nanowire Networks. In *2020 International Joint Conference on Neural Networks (IJCNN)*. 1–8. <https://doi.org/10.1109/IJCNN48605.2020.9207727>
- [13] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. 2009. The elements of statistical learning: data mining, inference, and prediction.
- [14] Sepp Hochreiter, A. Steven Younger, and Peter R. Conwell. 2001. Learning to Learn Using Gradient Descent. In *Artificial Neural Networks – ICANN 2001 (Lecture Notes in Computer Science)*, Georg Dorffner, Horst Bischof, and Kurt Hornik (Eds.). Springer, Berlin, Heidelberg, 87–94. [https://doi.org/10.1007/3-540-44668-0\\_13](https://doi.org/10.1007/3-540-44668-0_13)
- [15] Joel Hochstetter, Ruomin Zhu, Alon Loeffler, Adrian Diaz-Alvarez, Tomonobu Nakayama, and Zdenka Kuncic. 2021. Avalanches and Edge-of-Chaos Learning in Neuromorphic Nanowire Networks. *Nature Communications* 12, 1 (Dec. 2021), 4008. <https://doi.org/10.1038/s41467-021-24260-z>
- [16] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. 2020. Meta-Learning in Neural Networks: A Survey. arXiv:2004.05439 [cs, stat]
- [17] Daniele Ielmini and H-S Philip Wong. 2018. In-memory computing with resistive switching devices. *Nature electronics* 1, 6 (2018), 333–343.
- [18] Herbert Jaeger. 2001. The “Echo State” Approach to Analysing and Training Recurrent Neural Networks – with an Erratum Note. (2001), 47.
- [19] Herbert Jaeger. 2002. A Tutorial on Training Recurrent Neural Networks, Covering BPPT, RTRL, EKF and the “Echo State Network” Approach. (Oct. 2002).
- [20] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. 1983. Optimization by Simulated Annealing. *Science* 220, 4598 (1983), 671–680. [jstor:1690046](https://doi.org/10.1126/science.1190046)
- [21] Z. Kuncic, O. Kavehei, R. Zhu, A. Loeffler, K. Fu, J. Hochstetter, M. Li, J. M. Shine, A. Diaz-Alvarez, A. Stieg, J. Gimzewski, and T. Nakayama. 2020. Neuromorphic Information Processing with Nanowire Networks. In *2020 IEEE International Symposium on Circuits and Systems (ISCAS)*. 1–5. <https://doi.org/10.1109/ISCAS45731.2020.9181034>
- [22] Zdenka Kuncic and Tomonobu Nakayama. 2021. Neuromorphic Nanowire Networks: Principles, Progress and Future Prospects for Neuro-Inspired Information Processing. *Advances in Physics: X* 6, 1 (Jan. 2021), 1894234. <https://doi.org/10.1080/23746149.2021.1894234>
- [23] Sam Lilak, Walt Woods, Kelsey Scharnhorst, Christopher Dunham, Christof Teuscher, Adam Z. Stieg, and James K. Gimzewski. 2021. Spoken Digit Classification by In-Materio Reservoir Computing With Neuromorphic Atomic Switch Networks. *Frontiers in Nanotechnology* 3 (2021). <https://doi.org/10.3389/fnano.2021.675792>
- [24] Alon Loeffler, Adrian Diaz-Alvarez, Ruomin Zhu, Natesh Ganesh, James M Shine, Tomonobu Nakayama, and Zdenka Kuncic. 2023. Neuromorphic Learning, Working Memory, and Metaplasticity in Nanowire Networks. *SCIENCE ADVANCES* (2023). <https://doi.org/10.1126/sciadv.adg3289>
- [25] Alon Loeffler, Ruomin Zhu, Joel Hochstetter, Adrian Diaz-Alvarez, Tomonobu Nakayama, James M. Shine, and Zdenka Kuncic. 2021. Modularity and Multitasking in Neuro-Memristive Reservoir Networks. *Neuromorphic Computing and Engineering* 1, 1 (Aug. 2021), 014003. <https://doi.org/10.1088/2634-4386/ac156f>
- [26] Alon Loeffler, Ruomin Zhu, Joel Hochstetter, Mike Li, Kaiwei Fu, Adrian Diaz-Alvarez, Tomonobu Nakayama, James M. Shine, and Zdenka Kuncic. 2020. Topological Properties of Neuromorphic Nanowire Networks. *Frontiers in Neuroscience* 14 (March 2020), 184. <https://doi.org/10.3389/fnins.2020.00184>
- [27] Mantas Lukoševičius and Herbert Jaeger. 2009. Reservoir Computing Approaches to Recurrent Neural Network Training. *Computer Science Review* 3, 3 (Aug. 2009), 127–149. <https://doi.org/10.1016/j.cosrev.2009.03.005>
- [28] Wolfgang Maass, Thomas Natschläger, and Henry Markram. 2002. Real-Time Computing Without Stable States: A New Framework for Neural Computation Based on Perturbations. *Neural Computation* 14, 11 (Nov. 2002), 2531–2560. <https://doi.org/10.1162/089976602760407955>
- [29] Gianluca Milano, Enrique Miranda, and Carlo Ricciardi. 2022. Connectome of Memristive Nanowire Networks through Graph Theory. *Neural Networks* 150 (June 2022), 137–148. <https://doi.org/10.1016/j.neunet.2022.02.022>
- [30] Gianluca Milano, Giacomo Pedretti, Kevin Montano, Saverio Ricci, Shahin Hashemkhani, Luca Boarino, Daniele Ielmini, and Carlo Ricciardi. 2021. In Materia Reservoir Computing with a Fully Memristive Architecture Based on Self-Organizing Nanowire Networks. *Nature Materials* (Oct. 2021). <https://doi.org/10.1038/s41563-021-01099-9>
- [31] Ingo Rechenberg. 1973. Evolutionsstrategie. *Optimierung technischer Systeme nach Prinzipien der biologischen Evolution* (1973).
- [32] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution Strategies as a Scalable Alternative to Reinforcement Learning. <https://doi.org/10.48550/arXiv.1703.03864> arXiv:arXiv:1703.03864
- [33] Catherine D. Schuman, J. Parker Mitchell, Robert M. Patton, Thomas E. Potok, and James S. Plank. 2020. Evolutionary Optimization for Neuromorphic Systems. In *Proceedings of the Neuro-inspired Computational Elements Workshop*. ACM, Heidelberg Germany, 1–9. <https://doi.org/10.1145/3381755.3381758>
- [34] Frank Sehne, Christian Osendorfer, Thomas Rückstieß, Alex Graves, Jan Peters, and Jürgen Schmidhuber. 2010. Parameter-Exploring Policy Gradients. *Neural Networks* 23, 4 (May 2010), 551–559. <https://doi.org/10.1016/j.neunet.2009.12.004>
- [35] Henry O. Sillin, Renato Aguilera, Hsien-Hang Shieh, Audrius V. Avizienis, Masakazu Aono, Adam Z. Stieg, and James K. Gimzewski. 2013. A Theoretical and Experimental Study of Neuromorphic Atomic Switch Networks for Reservoir Computing. *Nanotechnology* 24, 38 (Sept. 2013), 384004. <https://doi.org/10.1088/0957-4484/24/38/384004>
- [36] Adam Z. Stieg, Audrius V. Avizienis, Henry O. Sillin, Cristina Martin-Olmos, Masakazu Aono, and James K. Gimzewski. 2012. Emergent Criticality in Complex Turing B-Type Atomic Switch Networks. *Advanced Materials* 24, 2 (2012), 286–293. <https://doi.org/10.1002/adma.201103053>
- [37] Anand Subramoney, Franz Scherr, and Wolfgang Maass. 2021. Reservoirs Learn to Learn. 59–76. [https://doi.org/10.1007/978-981-13-1687-6\\_3](https://doi.org/10.1007/978-981-13-1687-6_3) arXiv:1909.07486 [cs]
- [38] Daan Wierstra, Tom Schaul, Jan Peters, and Juergen Schmidhuber. 2008. Natural Evolution Strategies. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*. IEEE, Hong Kong, China, 3381–3387. <https://doi.org/10.1109/CEC.2008.4631255>
- [39] Tim Zhang, Amirali Amirsoleimani, Mostafa Rahimi Azghadi, Jason K Eshraghian, Roman Genov, and Yu Xia. 2023. SSCAE: A Neuromorphic SNN Autoencoder for sc-RNA-seq Dimensionality Reduction. (2023).
- [40] Ruomin Zhu, Joel Hochstetter, Alon Loeffler, Adrian Diaz-Alvarez, Tomonobu Nakayama, Joseph T. Lizier, and Zdenka Kuncic. 2021. Information Dynamics in Neuromorphic Nanowire Networks. *Scientific Reports* 11, 1 (June 2021), 13047. <https://doi.org/10.1038/s41598-021-92170-7>
- [41] Ruomin Zhu, Joel Hochstetter, Alon Loeffler, Adrian Diaz-Alvarez, Adam Stieg, James Gimzewski, Tomonobu Nakayama, and Zdenka Kuncic. 2020. Harnessing Adaptive Dynamics in Neuro-Memristive Nanowire Networks for Transfer Learning. In *2020 International Conference on Rebooting Computing (ICRC)*. 102–106. <https://doi.org/10.1109/ICRC2020.2020.00007>
- [42] Ruomin Zhu, Alon Loeffler, Joel Hochstetter, Adrian Diaz-Alvarez, Tomonobu Nakayama, Adam Stieg, James Gimzewski, Joseph Lizier, and Zdenka Kuncic. 2021. MNIST Classification Using Neuromorphic Nanowire Networks. In *International Conference on Neuromorphic Systems 2021 (ICONS 2021)*. Association for Computing Machinery, New York, NY, USA, 1–4. <https://doi.org/10.1145/3477145.3477162>
- [43] Rui-Jie Zhu, Qihang Zhao, and Jason K Eshraghian. 2023. SpikeGPT: Generative Pre-trained Language Model with Spiking Neural Networks. *arXiv preprint arXiv:2302.13939* (2023).

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009